

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

FINJAN SOFTWARE, LTD., an Israel)	
corporation,)	
)	
Plaintiff,)	
)	
v.)	C.A. No. 06-369-GMS
)	
SECURE COMPUTING CORPORATION, a)	DEMAND FOR JURY TRIAL
Delaware corporation; CYBERGUARD)	
CORPORATION, a Delaware corporation,)	
WEBWASHER AG, a German corporation and)	
DOES 1 THROUGH 100,)	
)	
Defendants.)	

**ANSWER AND COUNTERCLAIMS OF DEFENDANTS
SECURE COMPUTING CORPORATION, CYBERGUARD
CORPORATION, AND WEBWASHER AG TO PLAINTIFF'S AMENDED
COMPLAINT FOR PATENT INFRINGEMENT**

Defendants Secure Computing Corporation, CyberGuard Corporation, and Webwasher AG (collectively referred to as "Secure Computing") for their Answer to Plaintiff's Amended Complaint For Patent Infringement ("Amended Complaint"), state as follows:

GENERAL DENIAL

Secure Computing denies each and every allegation, matter or thing contained in the Amended Complaint which is not expressly admitted, qualified or answered herein.

THE PARTIES

1. Secure Computing lacks sufficient knowledge to form a belief as to the truth or falsity of the allegations contained in paragraph 1 of the Amended Complaint and therefore denies the same.

2. Secure Computing admits the allegations of paragraph 2 of the Amended Complaint.

3. Secure Computing denies the allegations of paragraph 3 of the Amended Complaint.

4. Secure Computing denies the allegations of paragraph 4 of the Amended Complaint.

JURISDICTION AND VENUE

5. The allegations in paragraph 5 are legal conclusions and do not require a responsive pleading. To the extent a response is required, Secure Computing does not dispute that jurisdiction is proper under 28 U.S.C. §§ 1331 and 1338.

6. The allegations in paragraph 6 are legal conclusions and do not require a responsive pleading. To the extent a response is required, Secure Computing does not dispute that venue is proper under 28 U.S.C. §§ 1391(b) and (c) and/or under 28 U.S.C. § 1400(b) and that personal jurisdiction is proper over Defendants. Secure Computing denies that Defendants CyberGuard Corp. and Webwasher AG are Delaware corporations. Secure Computing denies that any of the Defendants have and continue to infringe, contributorily infringe and/or induce infringement of U.S. Patent No. 6,092,194, U.S. Patent No. 6,804,780, and U.S. Patent No. 7,058,822.

PLAINTIFF'S PATENT

7. Secure Computing admits that Exhibit A to the Amended Complaint is a copy of the '194 patent. Secure Computing denies any and all other allegations of paragraph 7 of the Amended Complaint.

8. Secure Computing admits that Exhibit B to the Amended Complaint is a copy of the '780 patent. Secure Computing denies any and all other allegations of paragraph 8 of the Amended Complaint.

9. Secure Computing admits that Exhibit C to the Amended Complaint is a copy of the '822 patent. Secure Computing denies any and all other allegations of paragraph 9 of the Amended Complaint.

10. Secure Computing denies the allegations of paragraph 10 of the Amended Complaint.

PATENT INFRINGEMENT

11. Secure Computing Corporation admits that it is in the business of developing and distributing network and systems security solutions to organizations. Secure Computing denies any and all other allegations of paragraph 11 of the Amended Complaint.

12. Secure Computing admits that Defendant CyberGuard was in the business of developing and distributing information security solutions, but denies that CyberGuard is currently engaging in business. Secure Computing denies any and all other allegations of paragraph 12 of the Amended Complaint.

13. Secure Computing admits that Defendant Webwasher AG was in the business of developing and distributing Internet and email content security and filtering solutions, but denies that Webwasher AG is currently engaging in business. Secure Computing denies any and all other allegations of paragraph 13 of the Amended Complaint.

FIRST CAUSE OF ACTION

(Infringement of the '194 Patent)

14. Secure Computing admits that in paragraph 14 of the Amended Complaint Plaintiff re-alleges the allegations set forth in paragraphs 1 through 13 in the Amended Complaint and in response to paragraph 14 of the Amended Complaint, Secure Computing restates and incorporates by reference its answers to paragraphs 1 through 13 and denies any allegation not expressly admitted.

15. Secure Computing denies the allegations of paragraph 15 of the Amended Complaint.

16. Secure Computing denies the allegations of paragraph 16 of the Amended Complaint.

17. Secure Computing denies the allegations of paragraph 17 of the Amended Complaint.

SECOND CAUSE OF ACTION

(Infringement of the '780 Patent)

18. Secure Computing admits that in paragraph 18 of the Amended Complaint Plaintiff re-alleges the allegations set forth in paragraphs 1 through 17 in the Amended Complaint and in response to paragraph 18 of the Amended Complaint, Secure Computing restates and incorporates by reference its answers to paragraphs 1 through 17 and denies any allegation not expressly admitted.

19. Secure Computing denies the allegations of paragraph 19 of the Amended Complaint.

20. Secure Computing denies the allegations of paragraph 20 of the Amended Complaint.

21. Secure Computing denies the allegations of paragraph 21 of the Amended Complaint.

THIRD CAUSE OF ACTION

(Infringement of the '822 Patent)

22. Secure Computing admits that in paragraph 22 of the Amended Complaint Plaintiff re-alleges the allegations set forth in paragraphs 1 through 21 in the Amended Complaint and in response to paragraph 22 of the Amended Complaint, Secure Computing

restates and incorporates by reference its answers to paragraphs 1 through 21 and denies any allegation not expressly admitted.

23. Secure Computing denies the allegations of paragraph 23 of the Amended Complaint.

24. Secure Computing denies the allegations of paragraph 24 of the Amended Complaint.

25. Secure Computing denies the allegations of paragraph 25 of the Amended Complaint.

PRAYER FOR RELIEF

WHEREFORE, Defendants deny that Plaintiff is entitled to any judgment or relief in Plaintiff's favor, including the relief in paragraphs A through E of the Prayer for Relief in Plaintiff's Amended Complaint.

AFFIRMATIVE DEFENSES

FIRST AFFIRMATIVE DEFENSE

26. Plaintiff fails to state a claim upon which relief can be granted.

SECOND AFFIRMATIVE DEFENSE

27. Plaintiff's claims are barred by the doctrines of waiver, estoppel, and laches.

THIRD AFFIRMATIVE DEFENSE

28. Plaintiff's asserted patents, United States Patent No. 6,092,194 ("the '194 patent"), United States Patent No. 6,804,780 ("the '780 patent"), and United States Patent No. 7,058,822 ("the '822 patent"), are invalid and/or void for failure to satisfy the requirements of patentability contained in 35 U.S.C. Section 101, et seq., including, but not limited to, Sections 102, 103, and for lack of enablement and inadequate written description under 35 U.S.C. section 112, and for failure to disclose best mode.

29. Secure Computing reserves its right to assert additional grounds for the invalidity or unenforceability of the '194 patent, the '780 patent, and the '822 patent if it discovers such grounds during the course of the litigation.

FOURTH AFFIRMATIVE DEFENSE

30. Secure Computing does not infringe, has not induced infringement, and has not contributed to infringement of any valid and enforceable claim of any patent owned by Plaintiff.

FIFTH AFFIRMATIVE DEFENSE

31. Plaintiff's request for damages arising from alleged infringement, whether direct or otherwise, is barred by Plaintiff's failure to plead marking, failure to mark and/or failure to require licensees to mark.

SIXTH AFFIRMATIVE DEFENSE

32. The '194 patent is unenforceable due to Plaintiff's inequitable conduct under 37 C.F.R. § 1.27 and 37 C.F.R § 1.56.

33. Finjan engaged in inequitable conduct by affirmatively misrepresenting and/or failing to disclose material information with an intent to deceive or mislead the PTO. Specifically, Secure Computing's inequitable conduct claims are two-fold: (1) Finjan improperly claimed small-entity status to the PTO; and (2) Finjan made material misrepresentations to the PTO regarding prior art during prosecution.

34. The '194 patent claims the benefit of the filing date of U.S. Patent Application 60/030,639.

35. In U.S. Patent Application 60/030,639, Finjan attached an Appendix describing Finjan's software, SurfinGate, as an embodiment of the alleged invention.

36. In 1998, Finjan entered into an agreement with Cisco Systems, which had over 500 employees, to bundle Finjan's SurfinGate software with Cisco's PIX firewall.

37. On October 27, 1999, Finjan made fee payments to the PTO for the claims in the '194 patent and affirmed that "[a] small entity statement was filed in the prior nonprovisional application and such status is still proper and desired."

38. On information and belief, Finjan knew on October 27, 1999 that small-entity status was no longer proper.

39. On March 27, 2001, in the prosecution of U.S. Patent No. 6,209,103, Finjan paid fees to the PTO as a large entity.

40. Finjan did not notify the PTO in relation to the '194 patent, of its loss of small-entity status until December 12, 2003.

41. Finjan avoided a previous rejection by the PTO examiner during examination based on prior art disclosed in U.S. Patent No. 5,623,600 (Ji). Finjan stated, in the prosecution history, that "Ji teaches gateway detection of viruses attached to executable files, and does not teach hostile Downloadable detection. As is well known in the art, a Downloadable is mobile code that is requested by an ongoing process, downloaded from a source computer to a destination computer for automatic execution. The programs or documents of Ji are not Downloadables."

42. The Ji specification actually teaches that "the apparatus of the present invention could also be included on a FTP server or a world wide web server for scanning files and messages as they are downloaded from the web."

43. On information and belief, Finjan knew that the Ji specification teaches hostile Downloadable detection.

SEVENTH AFFIRMATIVE DEFENSE

44. Plaintiff's claims are barred and/or limited under 28 U.S.C. § 1498.

EIGHTH AFFIRMATIVE DEFENSE

45. Plaintiff's claims that Defendants are inducing infringement and/or contributorily infringing are barred because the underlying alleged direct infringers are licensed and/or released.

NINTH AFFIRMATIVE DEFENSE

46. Plaintiff's claims are barred and/or limited by the doctrine of patent exhaustion.

COUNTERCLAIMS

47. Defendant Secure Computing Corporation ("Secure Computing"), for its Counterclaims against Plaintiff, states:

48. Defendant Secure Computing Corporation is a corporation organized and existing under the laws of the State of Delaware, with its corporate headquarters at 4810 Harwood Road, San Jose, California 95124.

49. On information and belief, Plaintiff Finjan Software Ltd. is a corporation organized and existing under the laws of Israel, with its principal place of business at Hamachsheve St. 1, New Industrial Area, Netanya, 42504, Israel. On information and belief, Plaintiff Finjan Software Ltd. is in the business of developing and distributing network and systems security solutions to organizations.

50. These Counterclaims arise under the Patent Laws of the United States, Title 35, United States Code. The jurisdiction of this Court is founded upon and arises under Title 28, United States Code sections 1331 and 1338(a), and under the Federal Declaratory Judgment Act, Title 28, sections 2201 and 2202. Personal jurisdiction over Plaintiff comports with the United

States Constitution and is proper because Plaintiff has filed suit in this Court against Secure Computing asserting infringement of claims of United States Patent No. 6,092,194, United States Patent No. 6,804,780, and United States Patent No. 7,058,822 and because Plaintiff has and continues to infringe, contributorily infringe and/or induce the infringement of United States Patent No. 7,185,361 and United States Patent No. 6,357,010 in this district. Venue within this judicial district is proper under Title 28, United States Code sections 1391(b), 1391(c) and/or 1400.

COUNT I

DECLARATORY JUDGMENTS OF INVALIDITY, UNENFORCEABILITY, AND NONINFRINGEMENT

51. This is an action for a declaratory judgment, together with such relief based thereon as may be necessary or proper, pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. Sections 2201 and 2202. There is an actual controversy between Plaintiff and Secure Computing arising under the United States patent laws, Title 35 of the United States Code.

52. This Counterclaim arises under the Patent Laws of the United States, Title 35, United States Code. The jurisdiction of this Court is founded upon and arises under Title 28, United States Code sections 1338(a), and under the Federal Declaratory Judgment Act, Title 28, sections 2201 and 2202. Personal jurisdiction over Plaintiff comports with the United States Constitution and is proper because Plaintiff has filed suit in this Court against Secure Computing asserting infringement of claims of United States Patent No. 6,092,194, United States Patent No. 6,804,780, and United States Patent No. 7,058,822. Venue within this judicial district is proper under Title 28, United States Code sections 1391(b), 1391(c) and/or 1400.

53. Plaintiff alleges that it is the owner of the United States Patent No. 6,092,194, United States Patent No. 6,804,780, and United States Patent No. 7,058,822. Plaintiff contends

that Secure Computing infringes claims of United States Patent No. 6,092,194, United States Patent No. 6,804,780, and United States Patent No. 7,058,822.

54. United States Patent No. 6,092,194 and each claim thereof is invalid, unenforceable and not infringed by Secure Computing. U.S. Patent No. 6,804,780 and each claim thereof is invalid and not infringed by Secure Computing. U.S. Patent No. 7,058,822 and each claim thereof is invalid and not infringed by Secure Computing.

55. Secure Computing has not infringed any valid claims of U.S. Patent No. 6,092,194, either directly, indirectly, contributorily or otherwise, and have not induced any others to infringe said patent. Secure Computing has not infringed any valid claims of U.S. Patent No. 6,804,780, either directly, indirectly, contributorily or otherwise, and have not induced any others to infringe said patent. Secure Computing has not infringed any valid claims of U.S. Patent No. 7,058,822, either directly, indirectly, contributorily or otherwise, and have not induced any others to infringe said patent.

56. With respect to the '194 patent, Finjan engaged in inequitable conduct by affirmatively misrepresenting and/or failing to disclose material information with an intent to deceive or mislead the PTO. Specifically, Secure Computing's inequitable conduct claims are two-fold: (1) Finjan improperly claimed small-entity status to the PTO; and (2) Finjan made material misrepresentations to the PTO regarding prior art during prosecution. Secure Computing incorporates its allegations set forth in paragraphs 32-43 as if fully restated in this counterclaim.

57. Secure Computing therefore seeks a declaration and finding by this Court that United States Patent No. 6,092,194 is invalid and unenforceable and that Secure Computing does

not infringe, directly, indirectly, contributorily or otherwise, and has not induced any others to infringe, any valid claims of United States Patent No. 6,092,194.

58. Secure Computing therefore seeks a declaration and finding by this Court that U.S. Patent No. 6,804,780 is invalid and that Secure Computing does not infringe, directly, indirectly, contributorily or otherwise, and has not induced any others to infringe, any valid claims of U.S. Patent No. 6,804,780.

59. Secure Computing therefore seeks a declaration and finding by this Court that U.S. Patent No. 7,058,822 is invalid and that Secure Computing does not infringe, directly, indirectly, contributorily or otherwise, and has not induced any others to infringe, any valid claims of U.S. Patent No. 7,058,822.

COUNT II

PATENT INFRINGEMENT

60. On February 27, 2007, United States Patent No. 7,185,361 (hereinafter referred to as the “‘361 Patent”), entitled “System, Method and Computer Program Product For Authenticating Users Using a Lightweight Directory Access Protocol (LDAP) Directory Server,” was duly and legally issued to Thomas D. Ashoff, Steve O. Chew, Jeffrey J. Graham, and Andrew J. Mullican. Secure Computing was assigned all ownership rights to the ‘361 Patent. A true and correct copy of the ‘361 Patent is attached as Exhibit A.

61. On March 12, 2002, United States Patent No. 6,357,010 (hereinafter referred to as the “‘010 Patent”), entitled “System and Method For Controlling Access To Documents Stored On an Internal Network,” was duly and legally issued to Richard R. Viets, David G. Motes, Paula Budig Greve, and Wayne W. Herberg. Secure Computing was assigned all ownership rights to the ‘010 Patent. A true and correct copy of the ‘010 Patent is attached as Exhibit B.

62. Plaintiff, Finjan Software Ltd., has directly, indirectly, contributorily, and/or by inducement infringed the '361 Patent in violation of 35 U.S.C. §§ 271(a), (b), (c), and/or (f), literally and/or by the doctrine of equivalents, in this District and elsewhere in the United States, by making, using, selling, offering for sale, and distributing products, including but not limited to Vital Security Internet IBox™, and will continue to do so unless enjoined by this Court.

63. Finjan's infringement of the '361 Patent has caused injury to Secure Computing, and will continue to do so unless enjoined by this Court, thereby entitling Secure Computing to all remedies available under the Patent Laws of the United States, including 35 U.S.C. § 281-85. The continued infringement will cause irreparable injury and damage to Secure Computing for which Secure Computing has no adequate remedy at law.

64. Finjan has directly, indirectly, contributorily, and/or by inducement infringed the '010 Patent in violation of 35 U.S.C. §§ 271(a), (b), (c), and/or (f), literally and/or by the doctrine of equivalents, in this District and elsewhere in the United States by making, using, selling, offering for sale, and distributing products, including but not limited to its Vital Security™ for Documents (aka Finjan Mirage), and will continue to do so unless enjoined by this Court.

65. Finjan's infringement of the '010 Patent has caused injury to Secure Computing, and will continue to do so unless enjoined by this Court, thereby entitling Secure Computing to all remedies available under the Patent Laws of the United States, including 35 U.S.C. § 281-85. The continued infringement will cause irreparable injury and damage to Secure Computing for which Secure Computing has no adequate remedy at law.

PRAYER FOR RELIEF

WHEREFORE, Defendants Secure Computing Corporation, CyberGuard Corporation, and Webwasher AG deny that Plaintiff is entitled to any relief for its Amended Complaint and

prays for judgment in Defendants' favor as prayed for in their Answer and Counterclaims and against Plaintiff as follows:

A. That Plaintiff take nothing by its Amended Complaint and its claims against Defendants Secure Computing Corporation, CyberGuard Corporation, and Webwasher AG be dismissed with prejudice;

B. That Defendants Secure Computing Corporation, CyberGuard Corporation, and Webwasher AG be found not to infringe, either directly, indirectly, contributorily, or otherwise, and be found not to have induced infringement of any valid claims of United States Patent No. 6,092,194, United States Patent No. 6,804,780, and United States Patent No. 7,058,822;

C. That United States Patent No. 6,092,194 be found invalid and unenforceable;

D. That United States Patent No. 6,804,780 be found invalid;

E. That United States Patent No. 7,058,822 be found invalid;

F. A finding that Plaintiff has directly, indirectly, contributorily, and/or by inducement, literally and/or by the doctrine of equivalents, infringed United States Patent No. 7,185,361 in violation of 35 U.S.C. §§ 271(a), (b), (c), and/or (f);

G. A finding that Plaintiff has directly, indirectly, contributorily, and/or by inducement, literally and/or by the doctrine of equivalents, infringed United States Patent No. 6,357,010 in violation of 35 U.S.C. §§ 271(a), (b), (c), and/or (f);

H. A permanent injunction enjoining Plaintiff and its respective agents, servants, officers, directors, employees, subsidiaries, affiliates, joint venturers, and all persons acting in concert with them, directly or indirectly, from infringing, inducing others to infringe, or contributing to the infringement of United States Patent No. 7,185,361;

I. A permanent injunction enjoining Plaintiff and its respective agents, servants, officers, directors, employees, subsidiaries, affiliates, joint venturers, and all persons acting in concert with them, directly or indirectly, from infringing, inducing others to infringe, or contributing to the infringement of United States Patent No. 6,357,010;

J. An order that Plaintiff account for and pay Secure Computing the damages to which it is entitled as a consequence of patent infringement;

K. An order that Secure Computing be awarded prejudgment interest and its costs, disbursements and attorneys' fees herein in accordance with 35 U.S.C. § 285;

L. For an accounting of Plaintiff's sales, profits, royalties, and damages owing to Secure Computing;

M. That Defendants Secure Computing Corporation, CyberGuard Corporation, and Webwasher AG be awarded costs, disbursements and attorneys' fees incurred in defending themselves in connection with the Amended Complaint; and

N. That Defendants Secure Computing Corporation, CyberGuard Corporation, and Webwasher AG be awarded such other and further relief as the Court deems just, equitable and proper.

DEMAND FOR A JURY TRIAL

Defendants Secure Computing Corporation, CyberGuard Corporation, and Webwasher AG hereby request a trial by jury, pursuant to Rule 38 of the Federal Rules of Civil Procedure, on all issues so triable.



OF COUNSEL:

Ronald J. Schutz
Jake M. Holdreith
Christopher A. Seidl
Trevor J. Foster
Robins, Kaplan, Miller & Ciresi L.L.P.
2800 LaSalle Plaza
800 LaSalle Avenue
Minneapolis, MN 55402
(612) 349-8500

Fredrick L. Cottrell, III (#2555)
cottrell@rlf.com
Gregory E. Stuhlman (#4765)
stuhlman@rlf.com
Richards, Layton & Finger
One Rodney Square
P.O. Box 551
Wilmington, DE 19899
(302) 651-7700

Attorneys for Defendants

Dated: April 20, 2007

CERTIFICATE OF SERVICE

I hereby certify that on April 20, 2007, I electronically filed the foregoing with the Clerk of Court using CM/ECF which will send notification of such filing(s) to the following and which has also been served as noted:

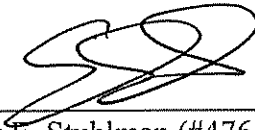
HAND DELIVERY

Philip A. Rovner
Potter Anderson & Corroon LLP
1313 N. Market Street,
Hercules Plaza, 6th Floor
P. O. Box 951
Wilmington, DE 19899-0951

I further certify that on April 20, 2007, the foregoing document was sent to the following non-registered participants in the manner indicated:

FEDERAL EXPRESS

Paul J. Andre
Perkins Coie LLP
101 Jefferson Street
Menlo Park, CA 94025-1114



Gregory E. Stuhlman (#4765)
stuhlman@rlf.com

EXHIBIT A



US007185361B1

(12) **United States Patent**
Ashoff et al.

(10) Patent No.: **US 7,185,361 B1**
(45) Date of Patent: **Feb. 27, 2007**

(54) **SYSTEM, METHOD AND COMPUTER PROGRAM PRODUCT FOR AUTHENTICATING USERS USING A LIGHTWEIGHT DIRECTORY ACCESS PROTOCOL (LDAP) DIRECTORY SERVER**

6,182,142 B1 * 1/2001 Win et al. 709/229
6,212,558 B1 * 4/2001 Antur et al. 709/221
6,233,688 B1 * 5/2001 Montenegro 713/201
6,324,648 B1 * 11/2001 Grantges, Jr. 713/201
2003/0126468 A1 * 7/2003 Markham 713/201

OTHER PUBLICATIONS

Microsoft Corporation, Microsoft Computer Dictionary, Microsoft Press, Third edition, p. 197.*
Definition of application gateway, Webopedia computer dictionary, http://www.webopedia.com/TERM/A/application_gateway.html *
Definition of firewall, Webopedia computer dictionary, <http://www.webopedia.com/TERM/f/firewall.html> *
Netegrity, SiteMinder 3.5 Architecture.
How to Securely Manage and Control User Access to E-Commerce Web Sites, Netegrity White Paper, Jul. 1999.
Check Point Account Management Client, Version 1.0, Sep. 1998.
FireWall-1 Architecture and Administration; Chapter 4, pp. 135-154, Sep. 1998.
Howes et al., *The LDAP Application Program Interface*, University of Michigan, Aug. 1995

* cited by examiner

Primary Examiner—Taghi T. Arani

(74) Attorney, Agent, or Firm—Schwegman, Lundberg, Woessner & Kluth, P.A.

(75) Inventors: Thomas D. Ashoff, Mt. Airy, MD (US);
Steve O. Chew, Pittsburgh, PA (US);
Jeffrey J. Graham, Olney, MD (US);
Andrew J. Mullican, Columbin, MD (US)

(73) Assignee: Secure Computing Corporation, St. Paul, MN (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days

(21) Appl. No.: 09/495,157

(22) Filed: Jan. 31, 2000

(51) Int. Cl.
H02H 3/05 (2006 01)

(52) U.S. Cl. 726/4; 713/151; 713/154;
707/1; 726/2; 726/8; 726/11; 726/12; 726/13;
726/14

(58) Field of Classification Search 713/201,
713/151-154; 707/1; 726/4, 8, 11-14
See application file for complete search history

(56) References Cited

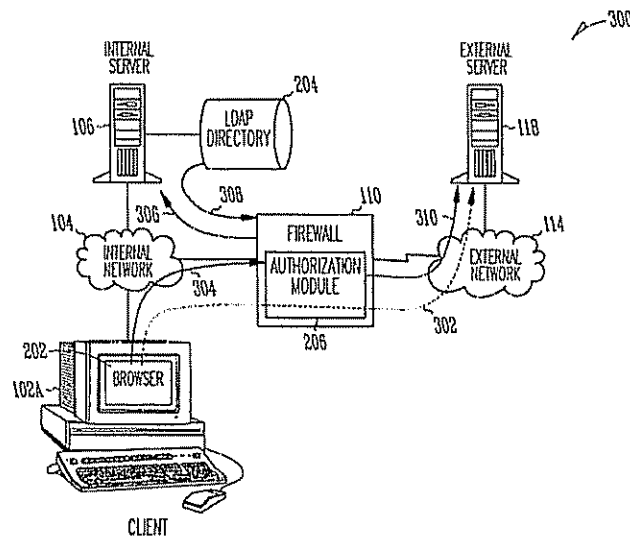
U.S. PATENT DOCUMENTS

5,657,390 A * 8/1997 Elgamal et al. 713/151
5,898,830 A * 4/1999 Wesinger, Jr. et al. 713/201
6,047,322 A * 4/2000 Vaid et al. 709/224
6,131,120 A * 10/2000 Reid 709/225

(57) ABSTRACT

A system, method and computer program product for providing authentication to a firewall using a lightweight directory access protocol (LDAP) directory server is disclosed. The firewall can be configured through a graphical user interface to implement an authentication scheme. The authentication scheme is based upon a determination of whether at least part of one or more LDAP entries satisfy an authorization filter.

15 Claims, 5 Drawing Sheets



U.S. Patent

Feb. 27, 2007

Sheet 1 of 5

US 7,185,361 B1

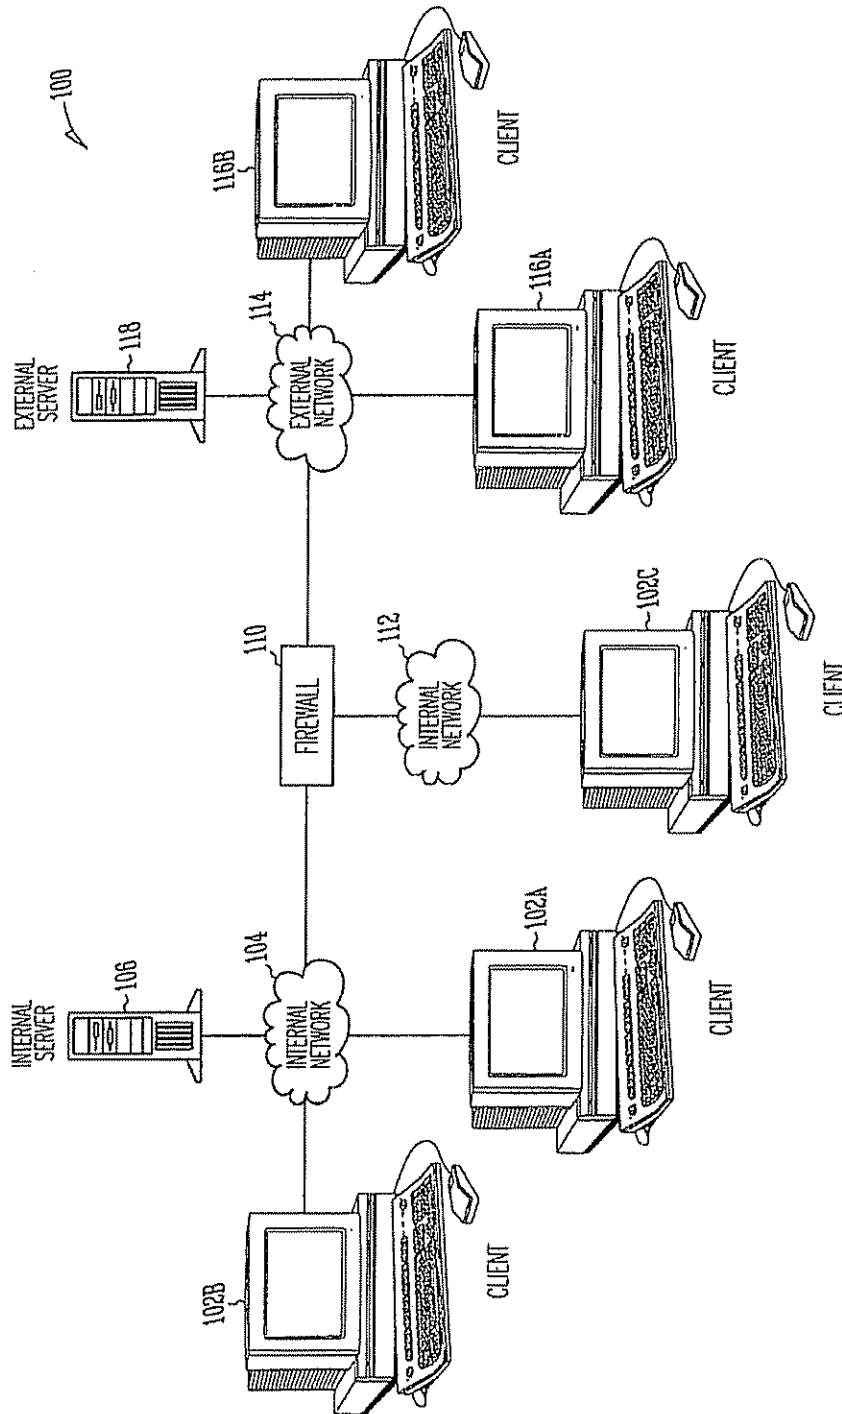
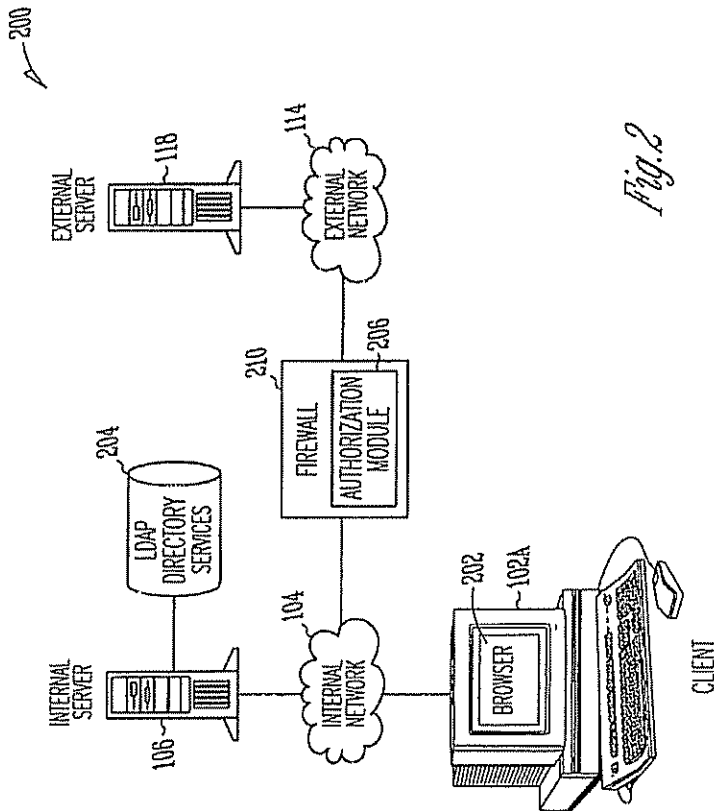


Fig. 1 (Prior Art)

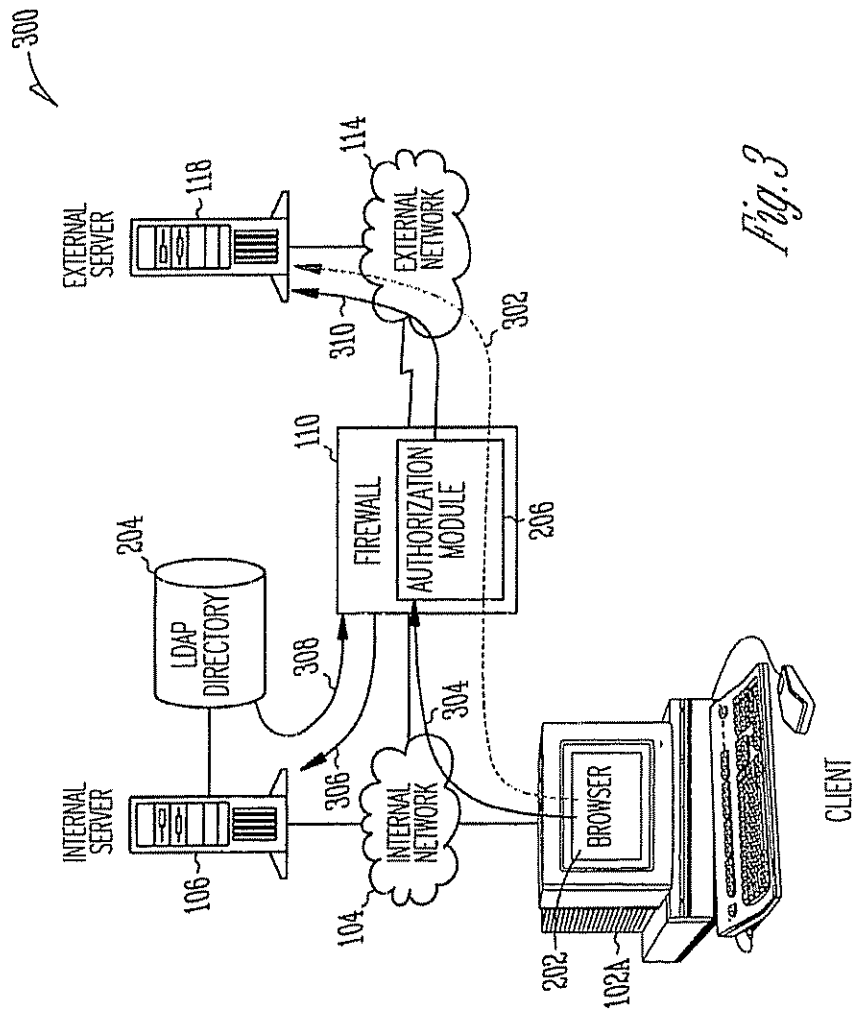


U.S. Patent

Feb. 27, 2007

Sheet 3 of 5

US 7,185,361 B1



U.S. Patent

Feb. 27, 2007

Sheet 4 of 5

US 7,185,361 B1

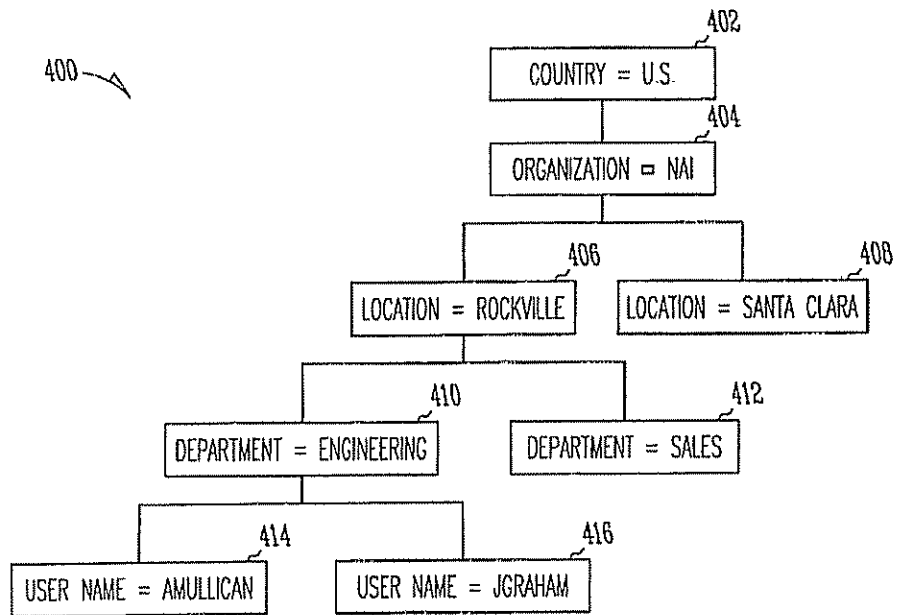


Fig. 4

U.S. Patent

Feb. 27, 2007

Sheet 5 of 5

US 7,185,361 B1

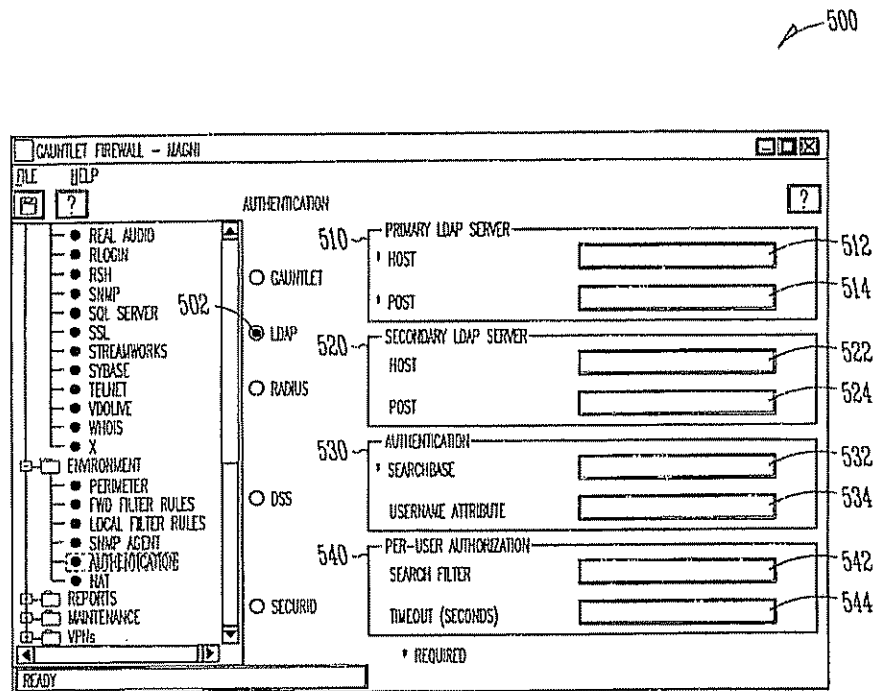


Fig. 5

US 7,185,361 B1

1

**SYSTEM, METHOD AND COMPUTER
PROGRAM PRODUCT FOR
AUTHENTICATING USERS USING A
LIGHTWEIGHT DIRECTORY ACCESS
PROTOCOL (LDAP) DIRECTORY SERVER**

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to user authentication mechanisms and more particularly to user authentication mechanisms for firewalls

2. Related Art

Control over access to information technology (IT) resources is a common need today. A firewall can be used to protect IT resources behind the firewall. Network firewalls can enforce a site's security policy by controlling the flow of traffic between two or more networks. For example, a company might encourage file transfers to the company's network that assist employees, but might discourage file transfers of potentially sensitive company confidential information from the company network to external destinations. Firewalls often are placed between a corporate network and an external network such as, e.g., the Internet, or a partner company's network. Firewalls can also be used to segment parts of a corporate network. A firewall system can provide both a perimeter defense to, e.g., an internal network, and a control point for monitoring access to and from specific networks such as, e.g., an external network.

Firewalls can control access at a network level, an application level, or both. At the network level, a firewall can restrict packet flow based on protocol attributes. For example, the packet's source address, destination address, originating transmission control protocol/user datagram protocol (TCP/UDP) port, destination port, and protocol type can be used for the control decisions. At an application level, a firewall can participate in communications between the source and destination applications with the firewall's control decisions being based on details of the conversation and other available information such as, e.g., previous connectivity or user identification. Thus, a firewall can authenticate users to control access to and from IT resources behind and before the firewall.

Firewalls can be packaged as system software, combined hardware and software, and, more recently, dedicated hardware appliances (e.g., embedded in routers, or easy-to-configure integrated hardware and software packages that can run on dedicated platforms). An example of an application-based firewall is the Gauntlet™ firewall available from Network Associates, Inc.

Firewalls can defend against attacks ranging from, e.g., unauthorized access, IP address "spoofing" (i.e., a technique by which hackers disguise traffic as coming from a trusted address to gain access to a protected network or resource), buffer overrun attacks, session hijacking, viruses and rogue applets, and rerouting of traffic. However, inherent limitations exist in certain services and protocols that conventional firewalls cannot remedy.

Conventionally, when software application programs sought to restrict what a user could do with the programs, the programs required identification of the user. For example, if a user desires access to sensitive corporate financial data in an accounting program, access to the data can be restricted by means of authentication mechanisms such as, e.g., a password. The application program therefore requires a list of users and identification information for the user for use in authenticating the user.

2

Early software application programs often included their own integrated authentication mechanisms. Users often use a variety of software application programs, each possibly having its own authentication mechanism. Users find it cumbersome to remember different passwords associated with each of the multiple software application programs.

IT resources used by companies today can include access to multiple software application programs and Internet based applications. For example, employees at a given company can use e-mail and groupware applications, and other office automation programs including, e.g., to spreadsheets, word-processors and presentation programs. As every application program conventionally has its own authentication mechanism, a separate database is initialized and updated for each application.

Authentication mechanisms can use a query to a database known as a directory that can store information about users. A directory is similar to a database in that one can store information in a directory and later retrieve the information from it. However, a directory is specialized in that a directory is typically designed for reading more than writing. A directory offers a static view of the information and allows simple updates without transactions. Thus, while a database is typically written to and read from frequently, a directory by comparison is primarily read from and is infrequently updated.

A directory service includes all the functions of a directory and adds a network protocol that can be used to access the directory. Standardization is desirable in implementing a directory service.

An early standard for directory service was the directory access protocol (DAP), which originated in the European standards organization. DAP although specifying a vast, feature-rich protocol for storing and encoding directory information, was unwieldy in size.

Today, a new protocol, lightweight directory access protocol (LDAP), is gaining wide acceptance in business. The LDAP standard defines an information model for a directory, a namespace for defining how directory information is referenced and organized, and a network protocol for accessing information in the directory. LDAP can also include an application programming interface (API). The LDAP protocol mandates how client and server computers can communicate with a LDAP directory. However, LDAP does not mandate how data should be stored. More and more companies today use an LDAP directory server to store a database of employees. The LDAP directory generally can store an employee name, phone number, address and other information about the employee, and a password for modifying the employee's information.

Firewalls also maintain a database of users and are operative to prompt users for an identifying user identifier and password. These conventional firewalls require that employee names and passwords be entered into a firewall authentication database. Maintenance of the firewall authentication database is especially burdensome where there are a large number of employees that are frequently leaving or joining a company or when a company has a large number of firewalls. Accordingly, what is needed is a mechanism for reducing this administrative burden. More specifically, what is needed is a mechanism for leveraging an existing LDAP directory server as part of a firewall's authentication process. In this manner, an existing LDAP directory server can be used as a central directory that stores the data used by all applications.

US 7,185,361 B1

3

SUMMARY OF THE INVENTION

A system, method and computer program product for enabling the authentication of users to a firewall using a lightweight directory access protocol (LDAP) directory server is provided by the present invention. The firewall can be configured through a graphical user interface to implement an authentication scheme. The authentication scheme is based upon a determination of whether information contained in one or more LDAP entries satisfy an authorization filter. It is a feature of the present invention that the authentication scheme can be configured independently of specifically stated field requirements or schema of the firewall. In accordance with the present invention, the authentication scheme can be flexibly specified to interact with a LDAP directory that has been uniquely developed for a company's internal needs. The company's investment in its existing administrative infrastructure can therefore be leveraged to a greater degree.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other features and advantages of the invention will be apparent from the following, more particular description of a preferred embodiment of the invention, as illustrated in the accompanying drawings.

FIG. 1 illustrates a communications network including a firewall.

FIG. 2 illustrates a communications network including a lightweight directory access protocol (LDAP) directory server and an authorization module within a firewall.

FIG. 3 illustrates the authentication of a client user through a firewall.

FIG. 4 illustrates an example embodiment of an LDAP directory tree.

FIG. 5 illustrates an embodiment of a graphical user interface for configuring the LDAP authentication feature.

DETAILED DESCRIPTION OF THE INVENTION

A preferred embodiment of the invention is discussed in detail below. While specific implementations are discussed, it should be understood that this is done for illustration purposes only. A person skilled in the relevant art will recognize that other components and configurations may be used without parting from the spirit and scope of the invention.

FIG. 1 illustrates an example embodiment of a communications network 100 including client computers 102a and 102b coupled via an internal network 104 to an internal server computer 106 and a firewall 110. Communications network 100 also includes a client computer 102c coupled via an internal network 112 to firewall 110. Finally, communications network 100 includes client computers 116a and 116b coupled via an external network 114 to an external server computer 118 and firewall 110. External network 114 can represent, e.g., the Global Internet, or a partnering company's network.

Network firewall 110 can enforce a business' security policy by controlling the flow of traffic between two or more networks such as, e.g., internal networks 104 and 112 and external network 114. In general, firewall 110 serves to isolate internal networks 104 and 112 from one another and also from external network 114.

As illustrated in FIG. 1, firewall 110 can be used to segment parts of a corporate network. For example, firewall

4

110 can be used to control information flow between a corporation's internal networks 104, 112. Firewall 110 can also provide a perimeter defense between an internal network 104, 112 and an external network 114.

FIG. 2 illustrates an example embodiment of a communications network 200 that includes client computer 102a coupled via internal network 104 to internal server 106 and to firewall 210. Firewall 210 is also coupled via external network 114 to external server 118.

As shown, client computer 102a includes a browser 202. Browser 202 can in one embodiment be an Internet browser that provides a graphical user interface to network resources. Browser 202 is generally operative to parse and make requests to network resources such as, e.g., external server 118, and present the results of the request to a client user viewing client computer 102a.

Internal server 106 is shown including a lightweight directory access protocol (LDAP) directory 204, which can be configured to store employee information. For example, a human resources database could be stored as an LDAP directory having a directory structure such as that illustrated in FIG. 4. As illustrated, LDAP directory tree 400 includes country 402 set in this example to US, organization 404 set to NAI, location 406 set to Rockville and location 408 set to Santa Clara, department 410 set to engineering and department 412 set to sales, and username 414 set to amullican and username 416 set to jgraham.

External server 118 can include an Internet server application. In one embodiment, the Internet server application supports file transfer protocol (FTP) communication. As would be apparent to those skilled in the relevant art, other types of server applications can be included on external server 118 including, e.g., databases, and electronic mail.

Firewall 210 is shown including an authorization module 206. Authorization module 206 is used to authenticate a client user (e.g., client computer 102a) to determine if the client user's communication is authorized to pass through firewall 210. Conventional firewalls 110 included their own database having a list of users and passwords, to enable authentication through firewall 110.

In accordance with the present invention, firewall 210 does not authenticate users using its own database. Rather, firewall 210 authenticates users using information contained within LDAP directory 204. As will be described in greater detail below, firewall 210 can authenticate users through an authentication scheme that can be based upon the unique composition of an organization's LDAP directory 204.

It is a feature of the present invention that the authentication scheme of the present invention can operate independently of specifically stated field requirements or schema of the firewall 210. In other words, an organization's LDAP directory 204 need not be modified to conform to a schema imposed by the firewall 210. Moreover, resistance to such a modification will not result in the maintenance of multiple directories.

In accordance with the present invention, the authentication scheme can be flexibly specified to interact with an existing LDAP directory that has been uniquely developed for a organization's internal needs. This framework enables a firewall administrator to seamlessly integrate a firewall product into an existing administrative infrastructure. The organization's investment in the existing administrative infrastructure can therefore be leveraged to a greater degree.

FIG. 3 illustrates the authentication process that is implemented by firewall 210. In the illustrated example, firewall 210 authenticates a client user at client computer 102a running a browser 202 that is attempting to access an

US 7,185,361 B1

5

application or resource on external server 118. This access path is illustrated by path 302.

This authentication process begins when client computer 102a initiates a network resource request 304 from browser 202. The network resource request 304 is intercepted by firewall 210. Authorization module 206 within firewall 210 challenges the client user to identify himself or herself. A challenge could in one embodiment include a request for entry of a username and password. Upon receipt of the identification information, authorization module 206 searches an authentication database (not shown) to identify an authentication method (e.g., LDAP authentication). If no entry in the authentication database is found for the client user, then a default authentication method can be used. In the LDAP authentication process, authorization module 206 binds to LDAP directory 204 and uses the userPassword attribute for authentication.

After authorization module 206 authenticates the client user, authorization module 206 then determines whether the client user is authorized to have his access request fulfilled. The LDAP authorization process is illustrated as communications 306 and 308. Communications 306 and 308 are facilitated using the LDAP protocol and may utilize the secure sockets layer.

If per-user authorization is configured, authorization module 206 determines whether one or more attributes of the client user's LDAP entry satisfies an authorization filter. If the one or more attributes of the client user's LDAP entry does not satisfy the authorization filter, then authorization module 206 determines that the authorization fails. If the authorization filter is satisfied, then the client user's network resource request is allowed through firewall 210. This allowed connection is illustrated in FIG. 3 as path 310.

To support per-user authorization, an administrator configures an authorization filter to use when authenticating users. One or more attributes in the client user's LDAP directory entry and associated values can be selected for the authorization filter. Once configured, authorization module 206 can verify that the LDAP entry used in the bind call satisfies the authorization filter before allowing the user access to/through the firewall.

FIG. 5 illustrates an example embodiment of a graphical user interface (GUI) 500 of a firewall systems administrator application screen. As shown by a selected radio button, LDAP authentication 502 has been selected. GUI 500 includes a primary LDAP server settings area 510, a secondary LDAP server settings area 520, an authentication settings area 530, and a per-user authorization settings area 540.

The primary LDAP server settings area 510 includes a host field 512 and a port field 514. The host field 512 can be used to enter an IP address or host name of a primary LDAP server. The port field 514 can be used to enter the port to be used on the primary LDAP server.

The secondary LDAP server settings area 520 also includes a host field 522 and a port field 524. The host field 522 can be used to enter an IP address or host name of a secondary LDAP server. The port field 524 is used to enter the port to be used on the secondary LDAP server. Fields 522, 524 can be left blank if no secondary LDAP server is being used.

The authentication settings area 530, can include searchbase field 532 and a username attribute field 534. The searchbase field 532 can be used to indicate the top of the directory tree 400 such as, e.g., country 402, organization 404, location 406, and department 410, so that a lookup can be within that portion of the directory tree. For example, a

6

set of attribute pairs such as, e.g., o=NAI, c=US to append to all requests to the LDAP server can be entered. The username attribute field 534 can include a default username attribute such as, e.g., uid. The username attribute field 534 can be used in performing per-user authorization.

The per-user authorization settings area 540 includes a search filter field 542 and a timeout field 544. The timeout field 544 can include a default value such as, e.g., 60 seconds. For example, timeout field 544 can be used to limit the amount of time the authorization filter query can take. If the time is exceeded, the authorization fails.

The search filter field 542 is used by firewall 210 in identifying the appropriate fields that are the subject of the LDAP directory authentication query. Upon receipt of a response from the LDAP directory 204, firewall 210 can then determine whether the client user is authorized to authenticate through the firewall 210.

In general, the authorization filter can contain any LDAP-valid combination of attributes and values, including object classes. At its simplest, the authorization filter specifies a single attribute and value pair. For example, the search filter field 542 can be used to enter a search filter expression such as "objectclass=gauntletUser".

Consider another example where LDAP directory 204 is configured by the company to include a field that would provide an access code level for each user. For example a "1" could correspond to only e-mail access, while a 5 could mean full access to all Internet services including world wide web browsing. In this environment, an authorization filter can be specified as "(&(objectclass=gUser)(status>=5))".

It should be noted that the authorization process need not be based on per-user authorization. In another embodiment, the authorization process can be based on a per-service authorization. In this embodiment, the per-service authorization can include an authorization for protocol services. Examples of protocol services include FTP, simple mail transport protocol (SMTP), e-mail, hypertext transport protocol (HTTP), etc. The per-service authorization can also be based on LDAP directory information. For example, authorization module 206 can use group memberships to determine whether a client user can use HTTP through firewall 210. To satisfy this authorization process, the authenticated user must be a member of the "web-users" group in the LDAP directory.

In one embodiment, the per-service authorization process uses the standard groupOfNames and groupOfUniqueNames object classes for authorization decisions. In general, a mechanism can be included that supports the specification of arbitrary group names for each service to be controlled. Control can then be based on a per-proxy basis or a per-policy basis.

Specification of per-service authorization criteria can also be implemented using the search filter field 542. In general, a different search (or authorization) filter can be provided for each service. For example, a search filter field can be included in GUI 500 to determine whether, e.g., a user is authorized to perform a file transfer, to send e-mail, or to access the world wide web. A search filter field can also be included in GUI 500 to determine whether, e.g., a user is a member of a particular group such as, e.g., engineering department 410, and if so, then particular services can be authorized based on being part of that group.

As noted, it is a feature of the present invention that firewall 210 can support arbitrary LDAP directory schema. Accordingly, firewall 210 does not require additional firewall-specific object classes or attributes in the directory

US 7,185,361 B1

7

Customers can populate the LDAP directories with whatever data they require. This authentication environment can be flexibly applied across multiple organizations each having their own sets of directory information. Indeed, the concepts of the present invention can be used to implement an authorization filter that relies on portions of information that are stored in distinct LDAP directories. This distributed authentication scheme enables an organization to implement segmented management of the user database.

While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

What is claimed is:

1. A system for authorizing client access to a network resource, comprising:

a server having at least one directory that can be accessed using a network protocol, said at least one directory being configured to store information concerning an entity's organization; and

a firewall that is configured to intercept network resource requests from a plurality of client users on an internal network, said firewall being operative to authorize a network resource request based upon a comparison of the contents of at least part of one or more entries in said at least one directory to an authorization filter, wherein said authorization filter is generated based on a directory schema that is predefined by said entity.

2. The system of claim 1, wherein said at least one directory is a lightweight directory access protocol directory.

3. The system of claim 1, wherein said authorization filter is specified using a graphical user interface.

4. The system of claim 1, wherein said authorization filter implements a per-user authentication scheme.

5. The system of claim 1, wherein said authorization filter implements a per-service authentication scheme.

6. The system of claim 1, wherein said firewall and said directory communicate using secure socket layer communication.

7. The system of claim 1, wherein said firewall is configured to query multiple directories.

8. An authentication method at a firewall, comprising the steps of:

(a) receiving a network resource request from a client user at an internal network;

(b) querying, using a network protocol, at least one directory that is configured to store information concerning an entity's organization, wherein said query is based upon an authorization filter that is generated based on a directory schema that is predefined by said entity;

8

(c) determining, based on the results of said query, whether the contents of at least part of one or more entries in said at least one directory satisfy said authorization filter; and

(d) permitting said network resource request through said firewall if said authorization filter is satisfied.

9. The method of claim 8, wherein step (b) comprises the step of querying said at least one directory using a lightweight directory access protocol.

10. The method of claim 8, further comprising the step of specifying an authorization filter using a graphical user interface.

11. The method of claim 10, wherein said specifying step comprises the step of specifying an authorization filter that implements a per-user authentication scheme.

12. The method of claim 10, wherein said specifying step comprises the step of specifying an authorization filter that implements a per-service authentication scheme.

13. The method of claim 8, wherein step (b) comprises the step of querying said directory using secure socket layer communication.

14. The method of claim 8, wherein step (b) comprises the step of querying multiple directories.

15. A computer program product for enabling a processor in a computer system to implement an authentication process, said computer program product comprising:

a computer usable medium having computer readable program code embodied in said medium for causing a program to execute on the computer system, said computer readable program code comprising:

first computer readable program code for enabling the computer system to receive a network resource request from a client user at an internal network;

second computer readable program code for enabling the computer system to query, using a network protocol, at least one directory that is configured to store information concerning an entity's organization, wherein said query is based upon an authorization filter that is generated based on a directory schema that is predefined by said entity;

third computer readable program code for enabling the computer system to determine, based on the results of said query, whether the contents of at least part of one or more entries in said at least one directory satisfy said authorization filter; and

fourth computer readable program code for enabling the computer system to permit said network resource request through a firewall if said authorization filter is satisfied.

* * * * *

EXHIBIT B



US006357010B1

(12) **United States Patent**
Viets et al.

(10) Patent No.: **US 6,357,010 B1**
(45) Date of Patent: **Mar. 12, 2002**

(54) **SYSTEM AND METHOD FOR CONTROLLING ACCESS TO DOCUMENTS STORED ON AN INTERNAL NETWORK**

(75) Inventors: Richard R. Viets, Naples; David G. Motes, Bonita Springs, both of FL (US); Paula Budig Greve, St. Anthony; Wayne W. Herberg, Rush City, both of MN (US)

(73) Assignee: Secure Computing Corporation, Roseville, MN (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: 09/024,576

(22) Filed: Feb. 17, 1998

(51) Int. Cl.⁷ G06F 12/14; G06F 15/173; H04L 12/00; H04L 9/00

(52) U.S. Cl. 713/201; 709/225; 713/200

(58) Field of Search 713/201, 200; 709/225

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,956,615 A	5/1976	Anderson et al.	235/61-7 B
4,177,510 A	12/1979	Appell et al.	364/200
4,584,639 A	4/1986	Hardy	364/200
4,621,321 A	11/1986	Boebert et al.	364/200
4,701,840 A	10/1987	Boebert et al.	364/200
4,713,753 A	12/1987	Boebert et al.	364/200
4,914,568 A	4/1990	Kadosky et al.	364/200
5,124,984 A	6/1992	Engle	370/94.1
5,179,658 A *	1/1993	Izawa	345/508
5,204,812 A *	4/1993	Kasiraj et al.	707/9
5,272,754 A	12/1993	Boebert	380/25
5,276,735 A	1/1994	Boebert et al.	380/21
5,311,593 A	5/1994	Carmi	380/23
5,329,623 A	7/1994	Smith et al.	395/275
5,455,953 A *	10/1995	Russell	710/266
5,544,321 A *	8/1996	Theimer et al.	714/9

5,566,170 A	10/1996	Bakke et al.	370/60
5,586,260 A	12/1996	Hu	395/200.2
5,606,668 A	2/1997	Shwed	395/200.11
5,619,648 A	4/1997	Canale et al.	395/200.01

(List continued on next page)

FOREIGN PATENT DOCUMENTS

EP	0697662 A1	2/1996	G06F/12/14
EP	0 743 777 A2	11/1996	H04L/29/06
EP	0811939 A2	12/1997	G06F/17/30
WO	97/13340	4/1997	H04L/9/00
WO	97/16911	5/1997	H04L/29/06
WO	97/26731	7/1997	H04L/9/00

OTHER PUBLICATIONS

Yialelis et al. "Role-Based Security for Distributed Object Systems", IEEE Proceeding, 1996, pp. 80-85.*

Sandhu et al. "Role-Based Access Control Models", IEEE Computer, Feb. 1996, pp. 38-47.*

Tari et al. "Role-Based Access Control For Intranet Security", IEEE Internet Computing, 1997, pp. 24-34.*

(List continued on next page)

Primary Examiner—Thomas Lee

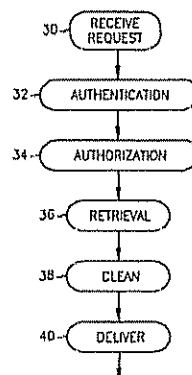
Assistant Examiner—Tanh Q. Nguyen

(74) Attorney, Agent, or Firm—Schwegman, Lundberg, Woessner & Kluth, P.A.

(57) **ABSTRACT**

A system and method of limiting access from an external network to documents stored on an internal network. A client list is built in which each client is assigned to one or more roles. Each role has access to one or more documents as defined on a document list. A request from an external network is reviewed and, if possible, the request is associated with a client on the client list. The requested document is then compared to the document list associated with the client's role and, if the requested document is in the list of documents available to a client in the client's role, the requested document is fetched, cleaned and sent to the client.

37 Claims, 6 Drawing Sheets



US 6,357,010 B1

Page 2

U.S. PATENT DOCUMENTS

5,623,601 A	4/1997	Vu	395/187.01
5,636,371 A	6/1997	Yu	395/500
5,673,322 A	9/1997	Pepe et al.	380/49
5,684,951 A	11/1997	Goldman et al.	395/188.01
5,689,566 A	11/1997	Nguyen	380/25
5,701,137 A *	12/1997	Kiernan et al.	345/340
5,708,780 A	1/1998	Levergood et al.	395/200.12
5,784,566 A *	7/1998	Viavant et al.	709/229
5,802,299 A *	9/1998	Logan et al.	709/218
5,819,271 A *	10/1998	Mahoney et al.	707/9
5,826,029 A *	10/1998	Gore, Jr. et al.	709/227
5,864,683 A *	1/1999	Boebert et al.	709/249
5,864,871 A *	1/1999	Kitain et al.	707/104
5,870,544 A *	2/1999	Curtis	713/201
5,884,033 A *	3/1999	Duvall et al.	709/206
5,884,312 A *	3/1999	Dustan et al.	707/10
5,892,905 A *	4/1999	Brandt et al.	713/201
5,903,732 A *	5/1999	Reed et al.	709/229
5,911,143 A *	6/1999	Deinhart et al.	707/103
5,913,024 A *	6/1999	Green et al.	713/200
5,915,087 A *	6/1999	Hammond et al.	713/201
5,918,013 A *	6/1999	Mighdool et al.	709/217
5,933,600 A *	8/1999	Shieh et al.	709/219
5,950,195 A *	9/1999	Stockwell et al.	707/4
5,961,601 A *	10/1999	Iyengar	709/229
5,987,611 A *	11/1999	Freund	713/201
6,023,765 A *	2/2000	Kuhn	713/200
6,055,637 A *	4/2000	Hudson et al.	713/201
6,088,679 A *	7/2000	Barkley	705/8

OTHER PUBLICATIONS

- International Search Report, PCT Application No. PCT/US 95/12681, 8 p. (mailed Apr. 9, 1996)
- Ancilotti, P., et al., "Language Features for Access Control", *IEEE Transactions on Software Engineering*, SE-9, 16-25 (Jan. 1983).
- Atkinson, R., "IP Authentication Header", Network Working Group, Request For Comment No. 1826, <http://ds.internic.net/rfc/rfc1826.txt>, 9 p. (Aug. 1995).
- Atkinson, R., "IP Encapsulating Security Payload (ESP)", Network Working Group, Request For Comment No. 1827, <http://ds.internic.net/rfc/rfc1827.txt>, 12 p. (Aug. 1995).
- Atkinson, R., "Security Architecture for the Internet Protocol", Network Working Group, Request for Comment No. 1825, <http://ds.internic.net/rfc/rfc1825.txt>, 21 p. (Aug. 1995).
- Baclace, P.E., "Competitive Agents for Information Filtering", *Communications of the ACM*, 35, 50 (Dec. 1992).
- Badger, L., et al., "Practical Domain and Type Enforcement for UNIX", *Proceedings of the 1995 IEEE Symposium on Security and Privacy*, p. 66-77 (May 1995).
- Belkin, N.J., et al., "Information Filtering and Information Retrieval: Two Sides of the Same Coin?", *Communications of the ACM*, 35, 29-38 (Dec. 1992).
- Bellovin, S.M., et al., "Network Firewalls", *IEEE Communications Magazine*, 32, 50-57 (Sep. 1994).
- Bevier, W.R., et al., "Connection Policies and Controlled Interference", *Proceedings of the Eighth IEEE Computer Security Foundations Workshop*, Kenmare, Ireland, p. 167-176 (Jun. 13-15, 1995).
- Bowen, T.F., et al., "The Datacycle Architecture", *Communications of the ACM*, 35, 71-81 (Dec. 1992).
- Bryan, J., "Firewalls For Sale", *BYTE*, 99-100, 102, 104-105 (Apr. 1995).
- Cobb, S., "Establishing Firewall Policy", *IEEE*, 198-205 (1996).
- Foltz, P.W., et al., "Personalized Information Delivery: An Analysis of Information Filtering Methods", *Communications of the ACM*, 35, 51-60 (Dec. 1992).
- Gassman, B., "Internet Security, and Firewalls Protection on the Internet", *IEEE*, 93-107 (1996).
- Goldberg, D., et al., "Using Collaborative Filtering to Weave an Information Tapestry", *Communications of the ACM*, 35, 61-70 (DEC. 1992).
- Grampp, F.T. "UNIX Operating System Security", *AT&T Bell Laboratories Technical Journal*, 63, 1649-1672 (Oct 1984).
- Greenwald, M., et al., "Designing an Academic Firewall: Policy, Practice, and Experience with SURF", *IEEE*, 79-92 (1996).
- Haigh, J.T., et al., "Extending the Noninterference Version of MLS for SAT", *Proceedings of the 1986 IEEE Symposium on Security and Privacy*, Oakland, CA, p. 232-239 (Apr. 7-9, 1986).
- Karn, P., et al., "The ESP DES-CBC Transform", Network Working Group, Request for Comment No. 1829, <http://ds.internic.net/rfc/rfc1829.txt>, 9 p. (Aug. 1995).
- Kent, S.T., "Internet Privacy Enhanced Mail", *Communications of the ACM*, 36, 48-60 (Aug. 1993).
- Lampson, B.W., et al., "Dynamic Protection Structures", *AFIPS Conference Proceedings*, 35, 1969 Fall Joint Computer Conference, Las Vegas, NV, 27-38 (Nov. 18-20, 1969).
- Lee, K.C., et al., "A Framework for Controlling Cooperative Agents", *Computer*, 8-16 (Jul. 1993).
- Lodin, S.W., et al., "Firewalls Fend Off Invasions from the Net", *IEEE Spectrum*, 26-34 (Feb. 1998).
- Loeb, S., "Architecting Personalized Delivery of Multimedia Information", *Communications of the ACM*, 35, 39-48 (1992).
- Loeb, S., et al., "Information Filtering", *Communications of the ACM*, 35, 26-28 (Dec. 1992).
- Merenbloom, P., "Network 'Fire Walls' Safeguard LAN Data from Outside Intrusion", *Infoworld*, p. 69 & addnl page (Jul. 25, 1994).
- Metzger, P., et al., "IP Authentication using Keyed MD5", Network Working Group, Request for Comments No. 1828, <http://ds.internic.net/rfc/rfc1828.txt>, 5 p. (Aug. 1995).
- Obraczka, K., et al., "Internet Resource Discovery Services", *Computer*, 8-22 (Sep. 1993).
- Peterson, L.L., et al., In: *Computer Networks*, Morgan Kaufmann Publishers, Inc., San Francisco, CA, p. 218-221, 284-286 (1996).
- Press, L., "The Net: Progress and Opportunity", *Communications of the ACM*, 35, 21-25 (Dec. 1992).
- Schroeder, M.D., et al., "A Hardware Architecture for Implementing Protection Rings", *Communications of the ACM*, 15, 157-170 (Mar. 1972).
- Schwartz, M.F., "Internet Resource Discovery at the University of Colorado", *Computer*, 25-35 (Sep. 1993).
- Smith, R.E., "Constructing a High Assurance Mail Guard", Secure Computing Corporation (Appeared in the Proceedings of the National Computer Security Conference), 7 p. (1994).
- Smith, R.E., "Sidewinder: Defense in Depth Using Type Enforcement", *International Journal of Network Management*, p. 219-229 (Jul.-Aug. 1995).
- Stadnyk, I., et al., "Modeling User's Interests in Information Filters", *Communications of the ACM*, 35, 49-50 (Dec. 1992).

US 6,357,010 B1

Page 3

Stempel, S., "IpAccess—An Internet Service Access System for Firewall Installations", *IEEE*, 31–41 (1995).
Stevens, C., "Automating the Creation of Information Filters", *Communications of the ACM*, 35, 48 (Dec. 1992).
Thomsen, D., "Type Enforcement: The New Security Model", *SPIE*, 2617, 143–150 (1995).
Warrier, U.S., et al., "A Platform for Heterogeneous Interconnection Network Management", *IEEE Journal on Selected Areas in Communications*, 8, 119–126 (Jan. 1990).
White, L.J., et al., "A Firewall Concept for Both Control-Flow and Data-Flow in Regression Integration Testing", *IEEE*, 262–271 (1992).
Wolfe, A., "Honeywell Builds Hardware for Computer Security", *Electronics*, 14–15 (Sep. 2, 1985).

Boebert, W.E., et al., "Secure Ada Target: Issues, System Design, and Verification", *Proceedings of the Symposium on Security and Privacy*, Oakland, California, Oakland, California, pp. 59–66, (1985).

Boebert, W.E., et al., "Secure Computing: The Secure Ada Target Approach", *Sci. Honeyweller*, 6(2), 17 pages, (1985).

Kahan, J., "A capability based authorization model for the world-Wide Web", *Computer Networks and ISDN Systems*, pp. 1055–1064, (1995).

Vinter, S.T., et al., "Extended Discretionary Access Controls", *IEEE*, pp. 39–49, (1988).

* cited by examiner

U.S. Patent

Mar. 12, 2002

Sheet 1 of 6

US 6,357,010 B1

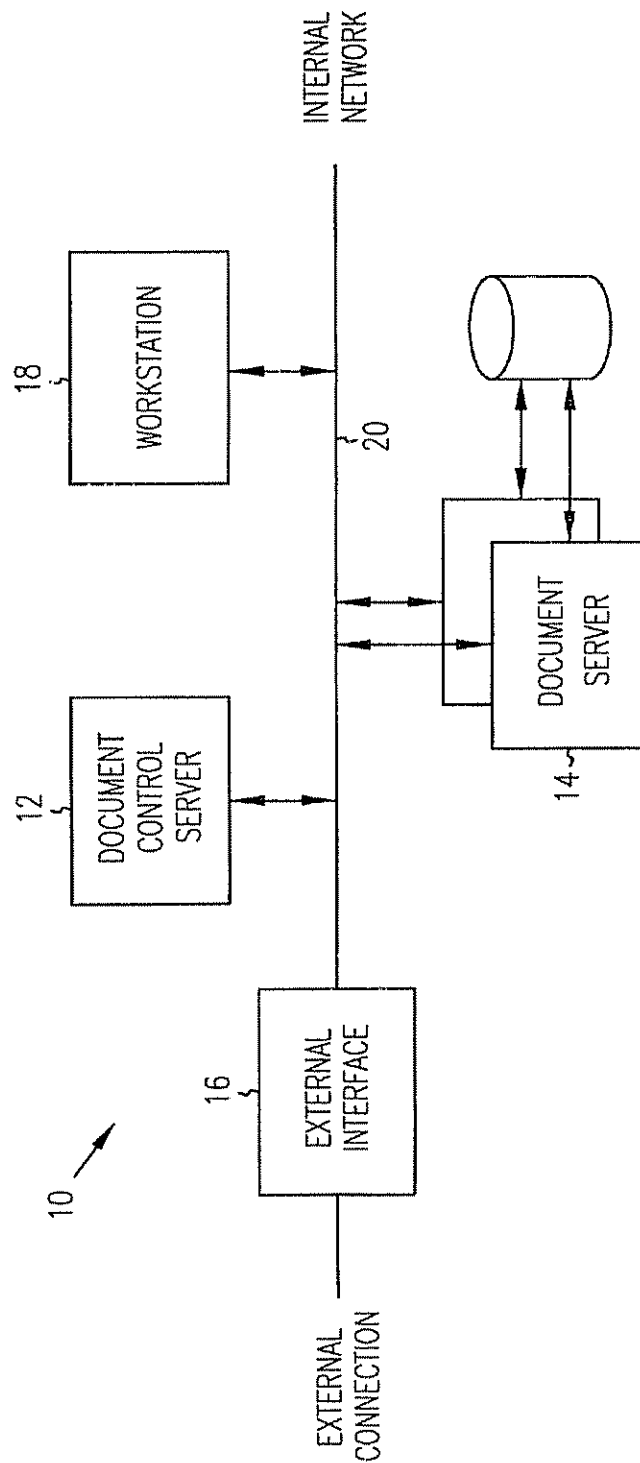


FIG. 1

U.S. Patent

Mar. 12, 2002

Sheet 2 of 6

US 6,357,010 B1

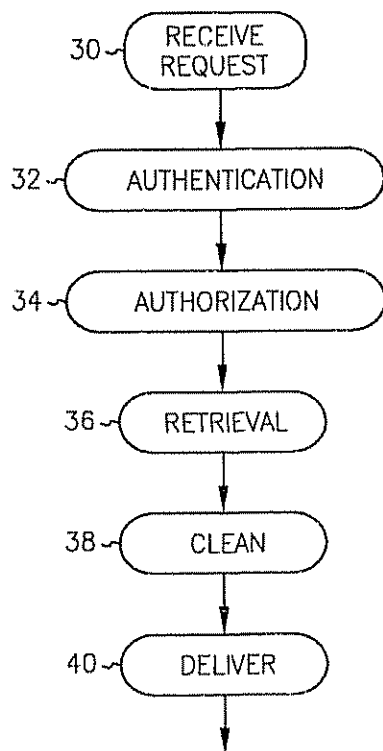


FIG. 2

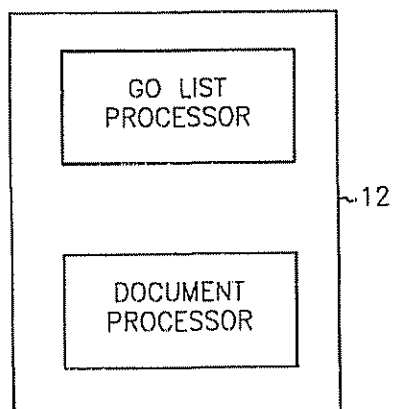


FIG. 3

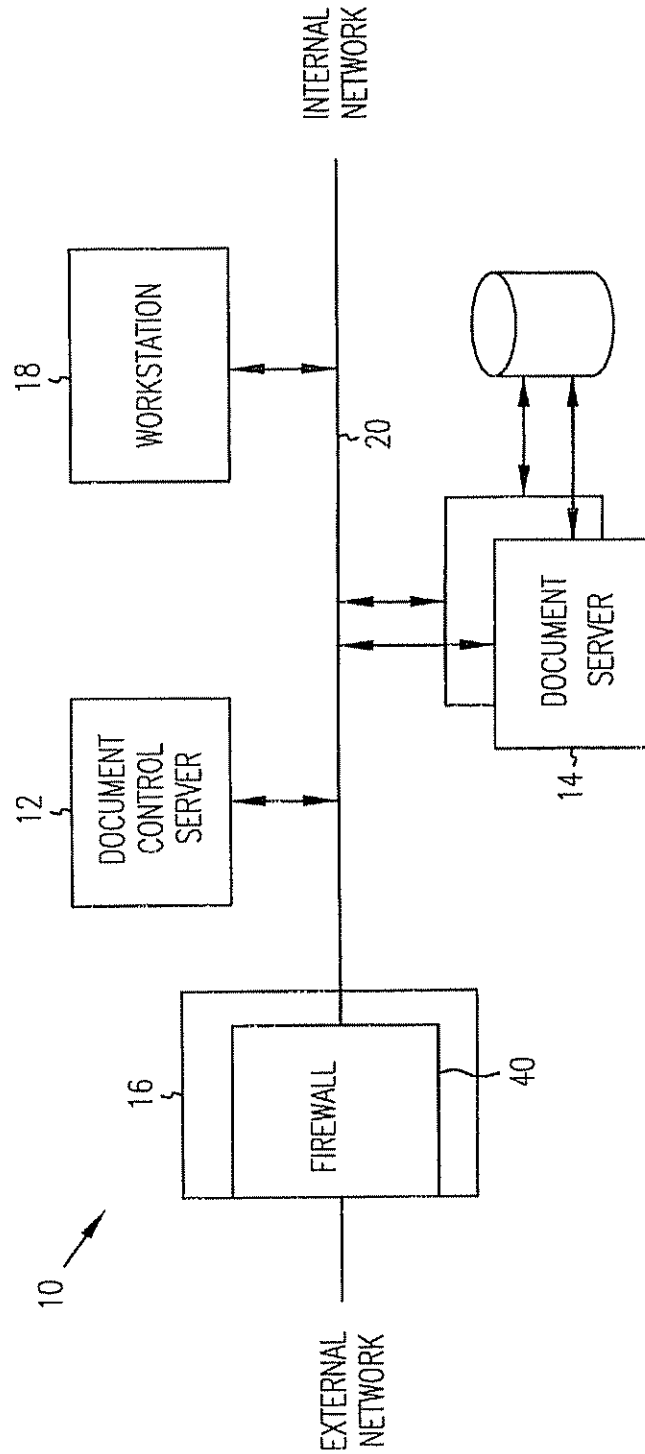


FIG. 4

U.S. Patent

Mar. 12, 2002

Sheet 4 of 6

US 6,357,010 B1

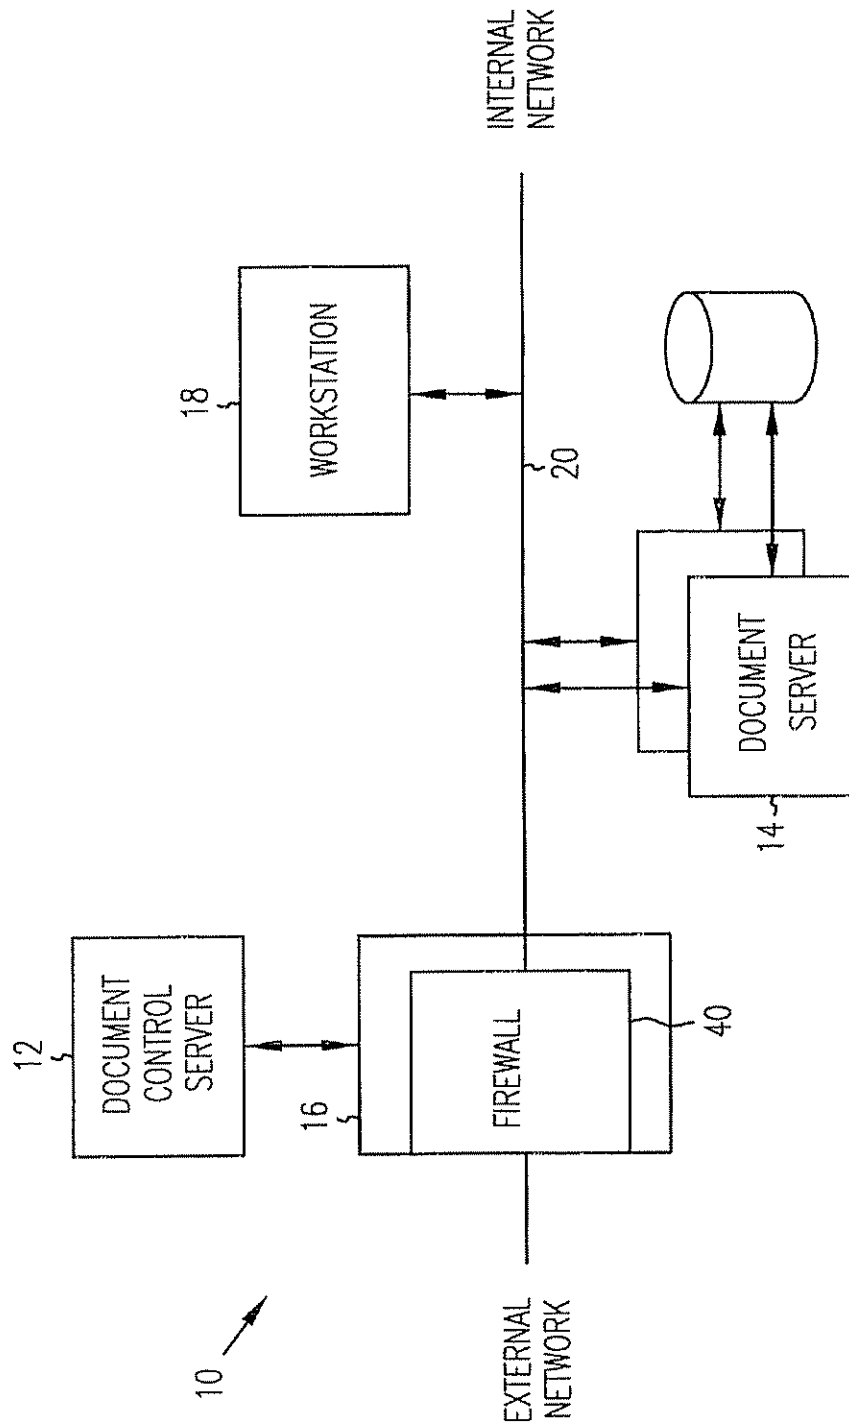


FIG. 5

U.S. Patent

Mar. 12, 2002

Sheet 5 of 6

US 6,357,010 B1

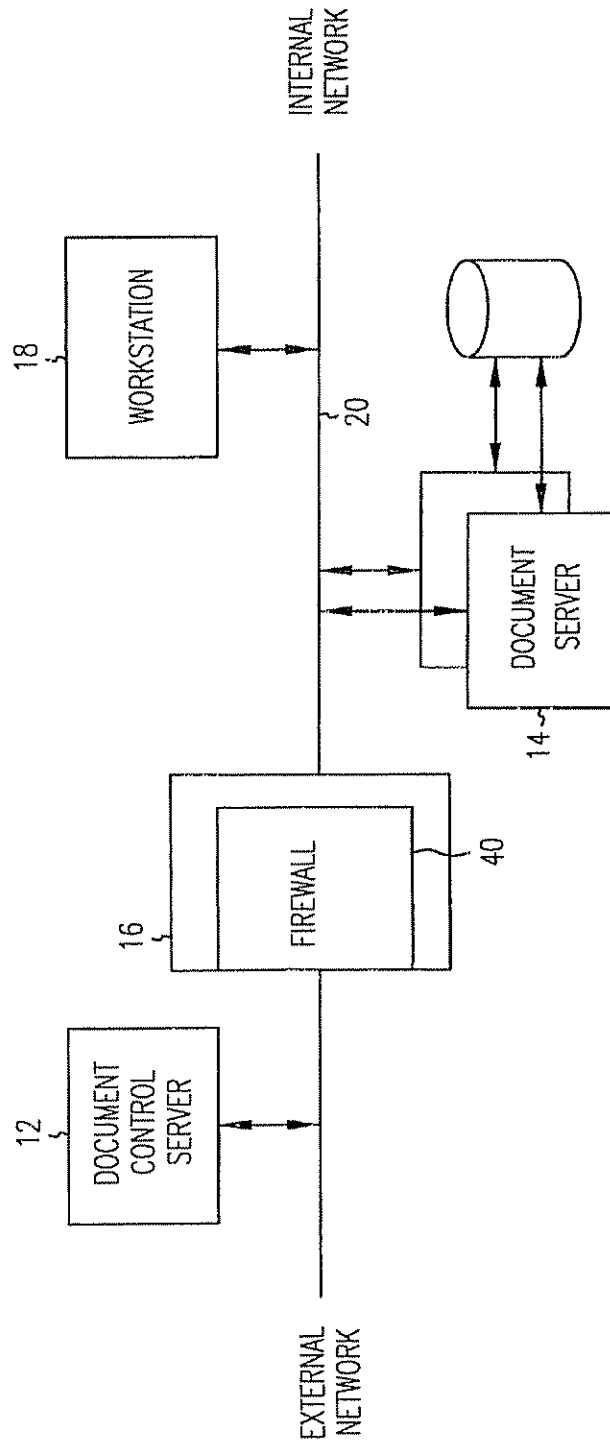


FIG. 6

U.S. Patent

Mar. 12, 2002

Sheet 6 of 6

US 6,357,010 B1

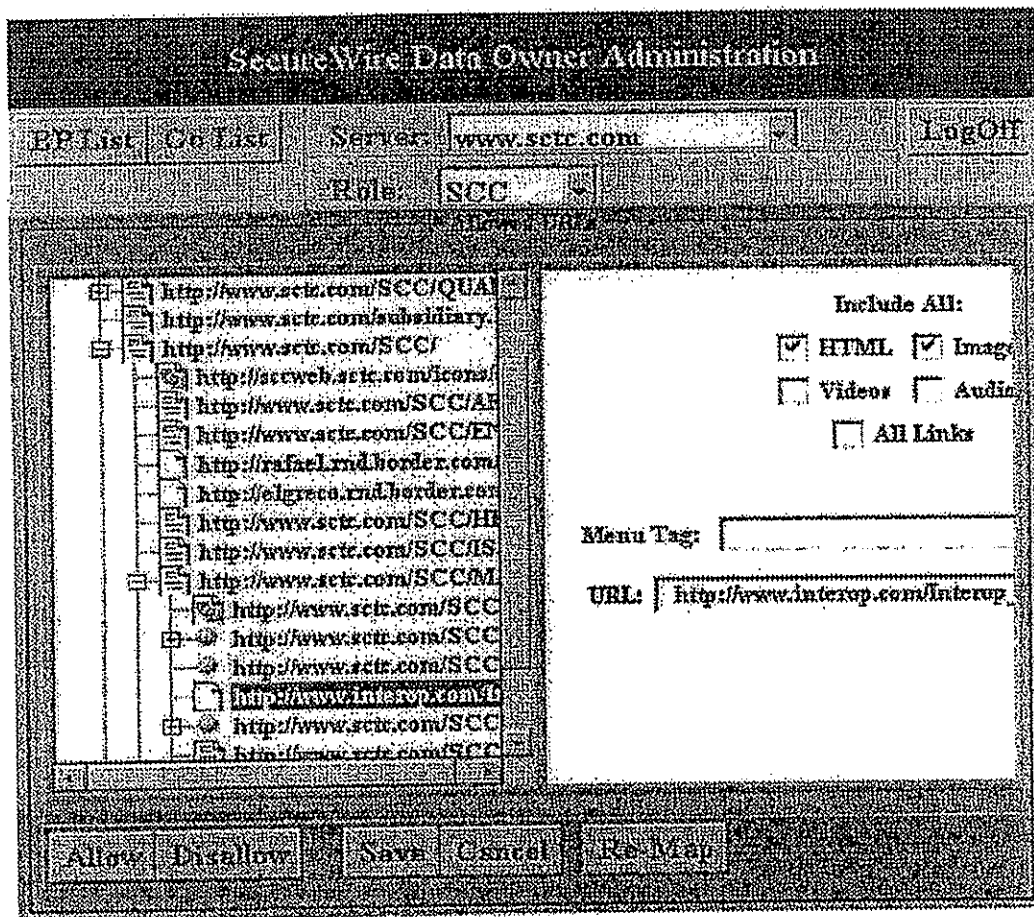


FIG. 7

US 6,357,010 B1

1

SYSTEM AND METHOD FOR CONTROLLING ACCESS TO DOCUMENTS STORED ON AN INTERNAL NETWORK

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to systems and methods for controlling communication between networks, and in particular to a system and method for limiting access to documents stored on an internal network

2. Background Information

Businesses today are acting cooperatively to achieve compatible business goals. For example, companies are using just-in-time manufacturing techniques to reduce overhead. To make this work, companies rely heavily on the ability of their suppliers to provide materials when needed.

At the same time, in this digital age business executives have become accustomed to receiving information from a number of sources both inside and outside the company almost instantaneously. They rely on such information to drive their day-to-day management decisions.

In order to provide outside organizations with relevant information in a timely manner, many companies have expanded their order-processing departments to handle increased call volumes. In this environment, outside partners call into the company's order-processing department to request specific information. This requires an employee to be available to answer calls, pull up information and verbally convey information to the partner. This option is very expensive, slow, and offers a poor level of service. What is needed is a system and method of streamlining the flow of information between partner companies while limiting access to company proprietary information.

The Internet provides one possible solution to this problem. The nature of the Internet makes it an ideal vehicle for organizations to communicate and share information. The Internet offers low cost universal access to information. Because of this, Internet transactions are expected to more than quadruple over the next two years, and partner communications via the Internet will almost double. Companies have begun to look to the Internet as a medium allowing quick, easy and inexpensive to business partners. To date, however, their Internet options have been limited.

One solution is to give business partners access to the company internal network. Companies are hesitant to do this, however, since such access, if abused, can lead to the disclosure of company sensitive information.

Another solution is to replicate necessary information to a web server located outside the company's firewall. Such an approach does allow organizations direct access to the information while at the same time limiting their access to company sensitive information. For this environment to work, however, the MIS department must manually transfer information from the internal network to the external server. Therefore, while this option offers organizations direct access to necessary data, that information can be 24 to 48 hours old. When dealing with just-in-time inventory levels and large dollar amounts, 24 hours is too late. This option also creates a bottleneck in MIS, redundancy of data, and decreased data integrity.

What is needed is a system and method for giving controlled access to designated documents stored on the internal network while restricting access to company sensitive information.

SUMMARY OF THE INVENTION

The present invention is a system and method of limiting access from an external network to documents stored on an

2

internal network. A client list is built in which each client is assigned to one or more roles. Each role has access to one or more documents as defined on a document list. A request from an external network is reviewed and, if possible, the request is associated with a client on the client list. The requested document is then compared to the document list associated with the client's role and, if the requested document is in the list of documents available to a client in the client's role, the requested document is fetched, cleaned and sent to the client.

According to another aspect of the present invention, a document control system is described. The document control system includes an internal network, an external interface, a document server connected to the internal network, and a document control server connected to the internal network and to the external interface. The document server controls access to a plurality of documents, including a first document. The document control server includes a go list processor for determining if the user has authorization to access said first document and a document processor for reading the first document from the document server, cleaning the first document and forwarding a clean version of said first document to the user. In operation, the document control server receives a document request from the external interface for the first document, determines a user associated with the document request, authenticates the user, determines if the user has authorization to access said first document and, if authorized, reads the first document from the document server, cleans the first document and forwards a clean version of said first document to the user.

BRIEF DESCRIPTION OF THE DRAWINGS

In the drawings, where like numerals refer to like components throughout the several views,

FIG. 1 shows a document access system;

FIG. 2 is a flow diagram illustrating operations performed by the document access system of FIG. 1;

FIG. 3 shows a document control server which can be used in the document access system shown in FIG. 1;

FIG. 4 is a document access system which includes a firewall;

FIG. 5 is a document access system in which the document control server is placed in a third network;

FIG. 6 is a document access system in which the document control server is placed on the external network; and

FIG. 7 is an example of a tree structure representation which could be used to aid the data owner in the selection of permitted URLs.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

In the following detailed description of the preferred embodiments, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. It is to be understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the present invention.

As noted above, corporations today are required by customers to deliver information such as price changes, new product data, manufacturing data, and customer support electronically. Competition is driving firms to work with partners through tight connections to internal systems. Allowing access, however, to such data in an efficient, manageable, and secure manner presents challenges. Com-

US 6,357,010 B1

3

panies go to great lengths to set up order processing departments and replicate large quantities of internal data to an external Internet server. These efforts are not only inefficient, but usually result in redundancy of data, decreased network integrity, and a bottleneck in the MIS department.

The present invention solves this problem by allowing specified external users controlled, customized, and secure access to the company's intranet without complex network infrastructure modifications. Further, the present invention permits one to control the parts of a Web server that are accessible to a Business Partner with only minimal intervention by IS personnel. (The term "Business Partner" is used in the following discussion to describe an external user who needs access to data such as Web pages which are not generally available to the public, but who also should not have unlimited to a company's intranet Web services.)

A document access system for giving controlled access to designated documents stored on the internal network while restricting access to company sensitive information is shown in FIG. 1. In FIG. 1 document access system 10 includes a document control server 12, a document server 14, an external interface 16 and one or more internal workstations 18. Document control server 12, document server 14, external interface 16 and internal workstations 18 are interconnected via internal network 20. Document server 14 reads and writes documents to storage 20. Requests for documents arrive at external interface 16 and are forwarded to document control server 12 for execution. In one embodiment external interface 16 includes a router used to form an Internet connection. In another embodiment, external interface 16 includes a direct connection interface such as formed by one or more modems used for direct dial-up by business partners wishing to access their data.

In one embodiment, as is illustrated in FIG. 2, at 30 document control server 12 receives a document request from the external interface for a first document. At 32, document control server 12 determines a user associated with the document request and authenticates the user. At 34, system 10 checks to see if the user is authorized to access the document requested. If so, at 36, system 10 retrieves the document from document server 14, cleans the document at 38 and, at 40, forwards the clean version of the document to the user. One embodiment of such a system and method is described next for a HyperText Transfer Protocol (HTTP) system.

When an HTTP or HTTPS connection request comes into document control server 12 there are three critical functions that must take place prior to returning the requested Web page: authentication, authorization, and internal connection. If either of the first two functions fail, the internal connection is not made. Then, once the internal connection has been made, document control server 12 must parse and "clean" the Web page prior to returning it to the requesting user.

Authentication

Authentication is fairly straight-forward and is of course visible to the end user. Following the HTTP protocol, when a user first enters a Uniform Resource Locator (URL) from their browser and the request is received at 30 (see FIG. 2) at server 12, a check is made at 32 for authentication. If basic authentication is being used, a check is made for authentication information in the HTTP header. If no username and password are found, the server returns a 401 error to the browser, telling the browser it needs to authenticate. The browser then pops up the box prompting the user to enter a username and password. When the HTTP request comes into the server, document control server 12 parses the username and password, comparing against its internal database of

4

users and if it finds a match, lets control proceed onto the authorization step. If no match is found, document control server 12 returns an error code back to the Web server it is running under (e.g., Internet Information Server or Netscape Enterprise Server) and the server will once again send back a 401 requesting the username and password. This process can happen 3 times and then the server will deny access. In one embodiment this authentication check is checked against a data base of known users as opposed to letting the Web Server check against a user database it may have.

Authorization

Once document control server 12 has authenticated the request, it must, at 34, determine if the user is authorized to get to the URL they have requested. This authorization will fail if the URL the user requested is not in the list of "allowed" URLs associated with the user. In one embodiment, each user is assigned one or more roles. Each role has access to a set of allowed URLs associated with that role.

In one embodiment each user has one or more roles associated with their user ID. For instance, they could be in the Marketing role, as well as the Engineering role. In one such embodiment, each role is directly associated with an internal server; you can only define one role for each server. This means you could not have the Marketing role and the Engineering role going to the same physical internal server. Such an approach can simplify system design.

In another embodiment, more than one roles may be assigned to each internal server. For example, a manufacturer may have all his reseller information on one server. One role, however, contains international resellers and another role contains domestic resellers. In such an embodiment, it would be advantageous to be able to define different sets of URLs on a single document server 14 that would allow for the different roles.

To complete the authorization portion, document control server 12 scans the list of allowed URLs for each role the user is in until it finds a match. If no match is made, an error condition is returned to the Web server indicating access is denied and the Web server in turn sends the appropriate error back to the browser.

An important and unique point to make here is that document control server 12 must translate the URL prior to doing its search for a match on the URL. When an external business partner (user) enters the URL, they enter a URL where the first part of the URL points document control server 12 and the second portion is the 'role' associated with that URL.

e.g., `https://www <server ID> com/Engineering/Standards/http_protocol.html`

where

`https=Secure http connection using SSL,`

`www <server ID> com=DNS name of document control server 12 (this name is unique to the customer installation)`

`/Engineering=role associated with this particular URL.`

`/Standards/http_protocol.html=actual web page on the internal web server`

If the real intranet web server associated with the Engineering role were `engineer.abcd.com`, then the translated URL that document control server 12 would search for is:

`engineer.abcd.com/Standards/http_protocol.html`

Note, the next two sections, Intranet connection and parsing the page, are entirely invisible to the end user.

Intranet Connection

If both the authentication and authorization phases completed successfully, document control server 12 will open a

US 6,357,010 B1

5

TCP connection to the appropriate intranet server (engineer.abcd.com from the above example) Once the TCP connection has been made, document control server 12 generates an http request for the specific web page. The intranet server locates and returns the requested web page to document control server 12.

Parsing the Page

The pages returned by the intranet are categorized as either text or non-text. Examples of the latter are graphics, such as GIF or JPEG documents, sound objects, or executable objects, such as Java applets. Non-text pages are not parsed and forwarded back to the client browser unchanged. Text documents, such as HTML formatted pages, however, contain embedded links that may need to be translated into their external equivalent. Embedded links fall into 3 groups, some of which require translation, while others don't: relative path links, server path links and absolute path links.

Relative path links, which are of the form subdir/page.html, don't require translation because the browser will prepend the path based on the referrer's page. For example, if the referrer's page was at:

`http://www.document_control_server.com/Engineering/Standards/http_protocol.htm`

and the relative link was `ssl_protocol.html`, then the browser would prepend

`http://www.document_control_server.com/Engineering/Standards/to the link.`

Server path links take the form of `/Specification/wheel.html` and require translation. This type of link points to a page that resides on the same server as the referrer's page, but with an absolute path starting at the root directory of the server. Assuming the same referrer's page as in the paragraph above, the translated link would be `/Engineering/Specification/wheel.html` (Note that the access string `http://` is not required because the browser will fill that in.) The translation is performed by prepending the alias associated with the referrer's page, Engineering, in this case, to the path of the embedded link.

Absolute path links are full URLs, such as

`http://engineer.abcd.com/Performance/testdrive.html`

and require translation only if they point to a server that is in document control server 12's Alias table. The example link will get translated because it points to the server engineer.abcd.com that exists in document control server 12's Alias table as Engineering. The translation is done by replacing the intranet server's name by the document control server 12's server name, followed by the alias of the intranet server. In this example, the translated URL would be

`http://www.<server 12 ID>.com/Engineering/Performance/testdrive.html`

Links that point to pages on servers unknown to document control server 12 are not translated because they may well point to valid external sites, such as Yahoo, which should be left untouched. In one embodiment, therefore, these links are not translated. (Note that if the referrer's page came in through the Secure Socket Layer (SSL), i.e., the URL starts with `https://`, then the translated links will also have `https://`.)

On the other hand, such links could pose a security threat. That is, the link could be pointing to an intranet server that contains sensitive information, whose existence should not be revealed to external users. To counter this, in one embodiment document control server 12 includes a list of links which should be hidden from the outside world. Links found on such a list would be translated to something innocuous. Redirection

When a page has moved, an intranet server may send a redirect status back to document control server 12. This

6

means that document control server 12 has to translate the redirected address, similar to how embedded links are handled, before forwarding it to the client browser.

Architecture

A document access system such as system 10 illustrated in FIG. 1 enables users to grant outside organizations direct access to internal web data in a secure, simple and manageable way. It is essentially a secure window through which outside partners can view internal web data. If, as is shown in FIG. 4, external interface 16 includes a firewall 40, system 10 also provides authenticated, authorized and view-customized business partner access to key Intranet servers via a standard web browser. Such a system enables users to easily, but accountably, grant authenticated partner access to internal web data, with complete control and authorization. Outside partners need only access a predefined URL in order to access an internal web page.

In one embodiment, as is shown in FIG. 4, document control server 12 is installed inside firewall 40. In another embodiment, as is shown in FIG. 5, document control server 12 is installed on a third network. Either way document control server 12 authenticates the outside user and then routes the request to a hidden, internal URL. The entire process is transparent to the outside user, and easily defined by the internal document control server 12 user. This allows business partners direct access to the data, eliminating the time lag, redundancy, lack of integrity and the bottleneck within the MIS Dept. (Please note that if document control server 12 is installed inside the firewall as is shown in FIG. 4, firewall 40 must be configured to restrict HTTP requests from any external source so that they can only get through to server 12. Similarly, if document control server 12 is installed on a third network (as a type of demilitarized zone (DMZ)) as is shown in FIG. 5, firewall 40 must be configured to restrict HTTP requests into internal network 20 so that they can only come from server 12.)

In a third embodiment, such as is shown in FIG. 6, document control server 12 is installed outside firewall 40 and is accessed through the Secure Sockets Layer (SSL). Such an embodiment should be set up so that firewall 40 allows HTTP traffic only from document control server 12 into internal network 20.

To further reduce any bottleneck, in one embodiment document control server 12 includes the option for the actual "data owners" themselves to define which partners have access to selective internal data. A data owner is a trusted individual within the organization that is empowered to grant Business Partners access privileges to Web pages on document servers 16. In one such embodiment, a Data Owner is assigned to one or more "roles," where a "role" represents the mapping alias assigned to one or the servers 16. A Data Owner can only add Business Partners or map URLs for the server "role" to which the Data Owner is assigned.

For example, an employee working in the Accounting department would be assigned to an Accounting role (server). The Accounting Data Owner is only able to access the internal servers specified by the administrator. This prevents the Accounting Data Owner from mapping URLs on any other server such as the Marketing or Engineering servers.

Once a Data Owner has been assigned to a role, he or she is able to perform the following tasks:

Add, modify, or delete a Business Partner from that particular role

Establish a user ID and password for a Business Partner for basic authentication

US 6,357,010 B1

7

Post or map an internal URL for access by a Business Partner

Delete URLs from a posted Go List

Delegation of such tasks to the data owners frees up MIS while also delegating data administration to those who understand the information best. In such an embodiment, the system administrator also defines general authentication rules and the list of eligible document servers 16.

Business Partners are somewhat-trusted end users. They can be granted controlled access to selected Web page structures on internal Web servers(s) such as document servers 16 once they have been provided with the following information:

A URL connecting them to document control server 12.

A user ID and password to authenticate them to document control server 12.

The name of the "menu tag" that they will select when they connect to document control server 12 that will retrieve the internal Web pages as specified by the Data Owner.

It is the Data Owners' responsibility to create and maintain a listing of Business Partners that require access to the intranet servers they control and provide the Business Partner with the information they will need to access the selected intranet server(s).

Every Business Partner defined by a Data Owner is part of a "group." The "group" a Business Partner belongs to is directly related to the role a Data Owner has been assigned and what internal servers are associated with that role. These groups control what URLs they are able to access on the internal servers.

A Business Partner can be assigned to multiple groups. For example, a Business Partner may belong to both the Marketing and the Sales groups. Data Owners manage their Business Partner accounts through the Business Partner list. From the Business Partner List, a Data Owner can establish a new Business Partner and modify or delete an existing Business Partner to any groups that they control.

In one embodiment, a Business Partner List is accessed by clicking on a BP List button on a Data Owner Administration utility window.

In one embodiment, document control server 12 includes a go list processor 22 and a document processor 24 (see FIG. 3). Go list processor 22 determines if the user has authorization to access said first document. Document processor 24 reads a document from document server 14, cleans the document as detailed above and forwards a clean version of the document to the user. Go list processor 22 and document processor 24 will be discussed next.

a) The Go List

The Go list is used by the document control server 12 to determine which URLs an authenticated Business Partner may be allowed to display. The Go List is unique to each role. It is identified by the rolename data within the roles/directory. In one embodiment, the Go List is managed by MIS. Such an embodiment does not, however, take advantage of the flexibility provided by the architecture of the present invention. Instead, it can be advantageous to permit individual data owners to determine the URLs to be included in each Go List. Such an embodiment will be discussed next. In this example, documents are made available by the Data Owner and can be accessed by a user termed a "business partner (BP)".

8

In one such embodiment, the Go list contains data formatted as follows:

Real_url; MENU="menu_name"

where the Real URL is the actual URL (without the http://) used by document control server 12 to access that particular directory or file. The MENU parameter is always present. There can be a value within the quotes, or it can be empty. If there is a value within the quotes, then document control server 12 will parse that value up and set up a link to that particular URL with the title of the link being the Menu Name—if the Business Partner goes to its menu page after it logs in.

Only allowed URLs are present within the go list. No other URLs are included.

Also, the Go List will permit the Business Partner to access any of the files under a certain directory. For the time being, this is done by default on any URL that the user allows that ends with a "/"—to allow the Business Partner to access anything within the subdirectory—this is entered in the go list with the traditional "*" following the trailing slash. In one embodiment, the directory URL is kept intact and the Data Owner is given the option to cut and paste the path that the Data Owner wants the whole directory included in. In such an embodiment Data Owners append the * to the directory name if they want everything within that directory accessible to the Business Partners within that role. In another embodiment, explicit "disallows" could be included to handle documents the Data Owner wants to except from inclusion in the list of accessible documents.

b) The Mapping Code on Document Control Server 12

The next portion of the whole mapping design is where a lot of the real work comes into play. There is some code within document control server 12 that gets called when the user (Data Owner) wants to map a particular server to the Go List. In one embodiment, a graphical user interface is used to select URLs and business partners. In one such embodiment, this code gets activated by the Data Owner performing one of the following tasks:

(A) Bringing up the Server Go List for the first time

(B) Clicking on a Node within the Go List Mapping Tree that has not yet been expanded and may have some children node

(C) Clicking on the Remap Button

At this point the GUI communicates to Server 12 via a Get URL request. The Get URL request:

(A) indicates to the server to load the GO list for a particular role

(B) checks to see if the node has not already been expanded and that the node exists on this server (the front part of the URL indicating the server name is consistent) and that the node is of an html type (or directory type)—if the node matches all of these criteria, then the GUI will indicate to the server to expand the particular URL for the particular role and

(C) If the DO selects the Remap button, they are prompted to see if they want to remap that portion of the server down. If the DO selects yes, then a request to Remap with the currently selected URL and the role is sent to the server.

Server 12 then acts upon the request that it receives from the GUI

(A) if the request was to load the GO list, then the server portion checks for the existence of a role_map data file in the roles directory. If this file does exist, then all that is done is that file is sent to the GUI line by line as is. If the file does not exist, then the file is created and the

US 6,357,010 B1

9

mapping function is called. The mapping function is called with the file pointer, the name of the URL (the server name) to be mapped, and a depth indicator of 1. (NOTE: This embodiment includes an option to go down multiple layers, however due to timeout issues it may be better to just go down one layer at a time and let the user build the map as they see fit. If it goes down multiple layers than the mapping function must have the skills and capabilities to prevent the same URL from being mapped multiple times (the recursive nature of links and web spiders). The mapping code and its behavior is described below these ordered steps.)

Once the mapping code is done, then the URLs with their appropriate line syntax have been input and saved into the mapping file. The mapping file is then sent line by line to the GUI.

(B) If the request was to expand a node, the server code then calls the mapping function for the particular URL to be expanded. It opens up a temp file to be used to write the information into. It then calls the mapping code with a file pointer to this temp file, the URL to be mapped, and a depth of 1. (NOTE: same code called as in condition A, just different parameters). Once the mapping code returns, the file is communicated line by line to the GUI and then the file is removed from the system.

(C) If the request was to Remap a particular server or URL, then things get a little tricky. If the DO had chosen to remap the entire server, then the URL sent is the base URL—otherwise the URL sent was the URL that the DO wanted to remap from that point down. The same code is used regardless of the situation—just a different URL value. What happens is:

The current map file is copied over to the new map file until the line with the URL to be remapped is read in. At this point this line is parsed to determine the depth in the tree (root level is 0).

A remap function is called which then does the following (note, that this is complicated by the fact that the tree could be at varying depths with files having been added or deleted and we would like to keep the prior shape and values of the tree when applicable).

Create a temp file to contain the intermediate results of the mapping.

Call the map function with a pointer to the temp file, the URL, and a depth of 1.

It then compares the current map file line by line with the temp file.

If the URL at the current depth is not found in the respective depth in the new temp file, that URL and any URLs immediately following it with a depth greater than the current depth are removed (that initial file is missing).

If the URL at the current depth is found in the temp file, any URL lines in between the URL being searched for and the prior URL are new files and are added prior to the current URL in the map file. Their syntax lines indicate the current depth and default of disallowing that URL. Then the existing current URL line is left as is in the map file.

At this point the next URL in the map file needs to be examined, in examining this URL the URL line syntax is examined. The depth is the issue of primary concern. If the depth is the same as the current depth, then continue with this loop. If the depth is a depth deeper

10

than the current depth, then this URL is a child of the prior URL and the prior URL then also needs to be remapped—the remap function is then called recursively with the prior URL. If the depth is less than the current depth, then we are no longer examining URLs that needed to be remapped and this function then returns.

After the final remap function returns, the remainder of the values are not to be touched and so they are copied over to the new map file as is.

After the server is done remapping, it then sends the whole newly mapped tree back to the GUI line by line. (NOTE: the server also then recreates the Go list to reflect the new values, information on creating the Go list from the mapping information is below.)

Finally, the GUI reads in the lines of information it gets from the server.

For (A) an initially loaded server, the GUI will create a tree examining each line of input. This is described in section 3 of this summary.

For (B) an expanded node, the GUI will create children nodes immediately below the expanded nodes by setting the depth according to the new tree and parsing each line of input. The parsing of the lines of input is described in section 3 of this summary.

For (C) a remapped server, the GUI will delete the previous tree and re-create the new tree (similar to A).

The code that does the mapping is a set of code pieces that loads the URL, parses any HTML that is returned, and creates a chain of information regarding the links. It then searches however deep is desired on each of the links.

It uses some standard libraries to aid in parsing and getting the URL and HTML.

when it encounters a link, it stores that information in memory. At that time it also attempts to determine what sort of file it is—currently we only distinguish between the following: HTML, sound, graphic, external, video. It attempts to determine this based off of the hints surrounding it based on the filename and the surrounding html.

If the file is an HTML file and the mapping code has not yet reached the requested depth of mapping, the mapping code with then recursively try to bring up the HTML file and parse through its contents, etc.

As it comes across a file and finishes parsing it—it creates a syntax line that the GUI is expecting and writes this line to the file that it is passed. The line is as follows:

URL (just a tag) http://real_file __url (the is the URL to the file that the mapping code downloaded and parsed—this is the file that will be allowed or denied by the DO)

depth (an integer that is to indicate the current depth that this file is in the tree—the lines are listed such that the tree can be loaded via a depth—first sort of algorithm—it continues down the left hand side while the depth is getting bigger, adding children from left to right as the depth is the same, and going back up the tree as the depth becomes smaller)

filename (this is the name of the file for that URL—a * is used if the filename cannot be determined (like in the case of a directory or the server))

file type (this is a character which corresponds to the filetypes that were previously mentioned)

state of node (this is a character which indicates if this node was in a collapsed or expanded state when the tree

US 6,357,010 B1

11

was saved—this is used by the GUI only, the mapping code always defaults this to C)
 allowed (this is a character which indicates if this URL is to be allowed or disallowed, the mapping code defaults this to disallowed)
 status of attaining the link (this is the HTTP status code of trying to access this link
 it could have a value of 200 which means it was accessed OK, 404 meaning that this URL was not found, or a 0 indicating that this was not accessed)
 already mapped (this is a one character flag used to indicate if this URL was already mapped and exists previously in the tree or not, this is most useful in a multi-depth mapping search)

Finally, the server has one more responsibility with respect to the mapping code. The server must handle writing out the saved data from the GUI when the DO requests to save the data. At this point, the data that is posted from the GUI to the server is written out line by line into the map file again. After this is done, the server deletes the prior Go List and parses through the map file a line at a time determining if that line is allowed. If the line is allowed the real_url part of the Map line is added (minus the http://) to the Go List. Next, the server checks the line to see if a "Menu" tag has been appended, if so, the server then adds the appropriate Menu tag to the Go list. Otherwise, it just adds a null menu tag to the Go list. As mentioned previously, currently if the real_url ends in a '/' an '*' is appended to the end of the real_url line to indicate that the user can access the entire directory—see Section 1 for further information.

After the Go List has been saved, the server is re-initialized with the new values—thus allowing immediate access or denial of access to the Business Partners for that role.

The Directory Map with the GUI

As mentioned previously, in one embodiment the GUI reads in and interprets the given line in order to creating nodes to represent the mapped server as a tree structure to the DO. Once such tree structure can be seen in FIG. 7.

The GUI communicates with server 12 as previously discussed and gets back well-defined lines of data. It parses the data and creates a node for each line provided by the Go list. It then looks at the expanded and collapsed feature to determine whether that node should be expanded or collapsed. It also looks at the file type to associate an image with that node. This image should allow the user to better determine what sort of nodes they are giving the Business Partners access to. The image is also determined by the return status from the mapping process—if a status of 404 is specified, then that link is determined to be broken (at least from document control server 12) and a broken link icon is displayed next to it. Finally, if that node is currently allowed, then the green ball is displayed next to it—to given the illusion of "green light means Go" to the DO.

The directory tree is created in a depth first fashion. It reads in each line examining the current depth. If the depth of the new node to be inserted is greater than the current depth than the node is inserted as a child of the previously inserted node. If the depth is the same, then it is inserted as a sibling. If the depth is less, then the tree is parsed back until the depth of the tree is at the same value of the depth and the node is inserted as a sibling at that level. The server is considered to be at the root level and thus has a depth of 0.

The text value that shows up with the node in the tree is the real_url value minus the http:// appended by the Menu Tag and Value if there is one for that node. A node may have a menu tag without being allowed.

12

The Data Owner (DO) can then traverse the tree examining the various links and determining what to allow and what to disallow. If they allow a directory—everything beneath that directory will be allowed. There currently is no mechanism for handling exceptions yet. If they allow files, then that file is allowed. If they have checked any other "include on allow" options (currently we offer, include all Gifs, Audio, Video, HTML, All Links)—then the files immediately beneath (once again only one layer deep) are automatically turned to allow if they are of the corresponding type. One note is that external files will never be "allowed"—as they do not exist on the server and thus it does not make sense for the DO to be allowing or disallowing those files.

When a data owner selects a link, they have the option of specifying the menu tag to be shown. If they do not specify one, then it is left blank. If they do then it is assigned for that node only. In order for it to get assigned to that node, the DO will have to select Allow or Disallow.

When Data Owners have finished making their changes they can save or cancel their mapping values. If they cancel then nothing that they did since their last save or remap will be saved. If they hit save, the values are communicated back to the server and the map file and the Go List File are updated as described previously.

Installation

Prior to install: Before fully installing and configuring document control server 12, the customer must have a digital certificate (for SSL encryption or transmissions) as well as set up and configure a server such as Internet Information Server (on NT) or Netscape Enterprise Server (on UNIX Solaris).

Installations: MIS installs document control server 12 onto the IIS or NES server and configures the firewall to allow HTTP access to the document control server 12 server.

Definition of end users: Document control server 12 offers organizations the option to delegate administration to end users who control the actual data (Data Owners) rather than forcing more work onto MIS. The data owner's access is defined by the network administrator. The data owner then maps which servers can be accessed. This data is stored in the document control server 12 "go list." The program code implementing document control server 12 is now completely installed, and ready for use.

Define business partner access: At this point, in order for an outside partner to access data, he/she must be granted access by the data owner. The data owner simply accesses the document control server 12 "data owner" GUI via a standard Java-enabled web browser. He/She can then define the new partner via role-based administration or explicitly choose which URLs may be accessed.

Outside users will not have access to any internal URL that is not specifically listed, even if there are embedded links in a URL for which access was granted. However, users can define access to a particular URL and all sub-pages as well.

Future partner access: After one role has been defined, future partners need only be added to that role, rather than requiring a whole new access to be defined.

Business partner access: All the outside partner has to do is type in the defined URL with any standard web browser. The partner will then be prompted for a user ID and password. Once these are entered, the partner will see a list of accessible URLs.

Back End Databases

It is important to understand that document control server 12 simply passes HTML information. This means document

US 6,357,010 B1

13

control server 12 does not have a problem passing CGI scripts and other dynamic content. Where this becomes particularly confusing is the access and authentication to back end databases via an application gateway

One of document control server 12's greatest values is to allow outsiders access to ever-changing information such as order processing, shipping, etc. Much of this data is stored in large back-end databases with an application gateway on the front end. The application gateway puts an HTML front end on the database and allows Intranet users to query required information. Typically, a user only needs to enter a customer account number to access this information. However, in order to give outside users direct access, many organizations need to require some level of authentication to this process.

When document control server 12 passes an outside partner to any Intranet URL, the user is authenticated as a unique document control server 12 user, however, that user ID is not passed on to the Intranet server. Therefore, direct access to back end databases cannot be defined for each document control server 12 Business Partner. It will be necessary to create an HTML-based sign-in screen. This is a simple process and offers an opportunity for resellers and professional services to add value to the product sale.

Many Internet web servers handle restricted access in slightly different ways. If the user is not known, the web server responds with a 401 error. The user's browser then displays a standard screen requesting user ID and password. The user types these in and is granted access.

It is important to understand that document control server 12 cannot process this transaction. For security reasons, only HTTP traffic can pass through document control server 12. Any authentication must be HTTP based, as mentioned above.

In one embodiment, document control server 12 is installed on a standard web server running IIS (NT) or NES (Solaris). It requires no changes to the current infrastructure, and no "agents" or "clients" to be installed on any web servers or browsers.

Administration and data owner usage is accessed via document control server 12's Java user interface. This allows access via any web browser that supports Java (e.g., Internet Explorer 4.0, Netscape 4.0). Outside partner access is also accomplished via a standard web browser.

Operating with Third Party Firewalls

Document control server 12 can be used in conjunction with a firewall to add an additional layer of security to Business Partner communications via the Web. Two components need to be considered when determining the location of document control server 12: Domain Name Service (DNS) and routing.

Depending on the deployment option preferred and the capabilities of firewall 40, there are up to four different methods for routing traffic to, and through, server 12:

1) Redirected proxy—For added security on external to internal connections, a redirected proxy can be configured on your firewall to redirect the inbound connection requests. When a Business Partner on the external network attempts to connect to document control server 12, firewall 40 intercepts the request and establishes a connection to server 12. This rerouted connection hides the actual destination from the Business Partner requesting the connection.

2) Transparent proxy—A transparent proxy can be set up through firewall 40 to document control server 12. From the Business Partners' perspective it will appear as though they are connecting directly to server 12 and not connecting to the firewall first.

14

3) Directly to document control server 12—If document control server 12 is installed on the external side of firewall 40 (as in FIG. 6), connection requests will be routed directly to server 12. In such an embodiment, server 12 authenticates the Business Partner, and passes the request through firewall 40. Firewall 40 then retrieves the requested Web page(s) from the specified document server 16.

4) Through a third network—(some firewalls allow a "third network" capability, (sometimes called the DMZ or the Secure Server Network). The three deployment scenarios discussed above still apply in a "three network" environment, however, additional firewall configuration is necessary to ensure that the required name resolution (DNS), and routing are still possible.

5) Security Features SSL encryption. In one embodiment, data transmitted between the partner and web server is SSL encrypted to prevent a sniffer from gathering information from the connection.

Document control server 12 server encrypted. Data stored on document control server 12 server such as user IDs, the go list, and partner profiles are all encrypted to prevent unauthorized access.

Password and user ID authentication. In one embodiment, document control server 12 supports password and user IDs for authorization. Stronger encryption could also be used. Granular Access Controls.

Business partners can only access internal URLs to which explicit access is given. If an accessible URL has embedded links to pages to which explicit access has not been granted, the partner cannot connect to them. However, if an embedded link is to an outside server, such as www.yahoo.com, in one embodiment access will not be restricted.

Internal URLs and IP Address Are Hidden

To ensure the security of the internal network and web pages, internal URLs and IP addresses are hidden from outside access. Partners type in a predefined URL and are presented with a list of accessible internal URLs. When a link is selected from the list Document control server 12 then maps to the internal URL. The internal URL and IP address are never displayed for the partner to see.

System Requirements, Compatibility and Performance

Considerations for performance and reliability include amount of cache memory, CPU power, BUS speed, amount of RAM, speed of memory chips, bus architecture (IDE, EIDE, PCI etc.), hard drive capacity, and hard drive quality (seek and access speeds). The following table identifies the basic characteristics of minimum, recommended and ideal server configurations to run document control server 12.

System Component	Minimum	Recommended	Ideal
CPU	Pentium 166	Pentium 200	Pentium Pro 200
RAM	32 MB	48 MB	64 MB
Hard disk	1 GB	2 GB	4 GB
Platform	IIS 3.0 (NT 4.0) or NES 3.0 (Solaris 2.5.1)		
Browser	Java enabled web browser MS Internet Explorer 4.0 or higher Netscape Navigator 4.0 or higher with Netscape's JDK 1.1 patch		
Other	CD-ROM, 3.5" diskette, Color monitor. Keyboard, Mouse		

Document control server 12 enables users to easily, but accountably, grant authenticated partner access to internal web data, with complete control and authorization. Outside partners need only access a predefined URL in order to access an internal web page.

US 6,357,010 B1

15

The following examples provide a better idea of how document control server 12 can be used to meet a variety of needs.

Example—Manufacturing (Order Processing)

Manufacturing companies process thousands of orders every day. In order to keep up with competition and to achieve the highest levels of quality, customers/partners need to know the immediate status of an order to the minute. Many companies today have moved to just-in-time inventory systems to reduce overhead and costs. Document control server 12 can grant access directly to an order-processing page that connects directly into an order-processing database. The order-processing agents (data owners) can define what data customers/partners have direct access to. As a result, the customer knows immediately the status of an order. The supplier also saves money by eliminating the need to replicate data or take phone calls asking for updates.

Example—Distribution

A distribution environment operates an order-processing and shipping department very similar to manufacturing. However, distribution also requires various types of information to be distributed to different partners, such as pricing and quantity breaks. Document control server 12 allows a company to customize the view each distributor or reseller sees, such as pricing or quantity breaks.

Example—Financial Services

Financial institutions process millions of transactions a day with a large number of outside partners. These include the purchase and sale of assets as well as order/sale confirmation, etc. Today, many of these transactions require a third party to set up a secure certificate. Document control server 12 can speed up this entire process by allowing an agent to immediately allow an outside customer or partner access to trading information in minutes, and without the need for third-party intervention.

Example—Health Services

Health care and insurance organizations process thousands of claims each day. Partners need a secure way to pass medical information and process it into a company's systems. For example, a doctor treats a patient who has Blue Cross/Blue Shield. That doctor needs to know if the patient's insurance covers the treatment, then process the claim after the treatment is given, and finally check on the status of payment once a claim is submitted. With document control server 12, Blue Cross/Blue Shield can give the doctor's office access to their internal list of insured patients, as well as the status of current claims. The company no longer needs to replicate this data to a DMZ or SSN Internet server or handle a phone call. The doctor's office can also securely fill out a web-based claim form over the Internet to process the claim for treatment.

Example—Government Agency

It is necessary for various government agencies and departments to frequently share sensitive data. One example is the CIA and various law enforcement agencies. The FBI, DEA, ATF and other agencies must routinely check into the files of various personnel and public citizens. Typically, this requires these agencies to send a paper request for information to the CIA. The CIA must then search for the relevant information and then send a copy back to the requesting agency.

With document control server 12, the FBI and other agencies can be given direct access to the CIA files that might be relevant such as histories and fingerprint analysis databases. This can save time and money.

16

Document control server 12 offers several advantages over current methods such as cost savings, improved customer service and leveraging of the current infrastructure. Current methods for passing data to outside partners are expensive, slow and unreliable. Document control server 12 offers the information to partners faster, easier and cheaper. It also more tightly integrates partners, thus improving business relations. Document control server 12 also leverages the benefits of current technology such as the Internet and Intranet.

Other business advantages of document control server 12 include: it reduces overhead and costs; it eliminates the need to copy content to a web server within the DMZ or external network; it offers spontaneous, dynamic user-managed content; it eliminates the wait for an IS manager to update data or post on a web server; it eliminates integrity and replication issues; it more tightly integrates partners; and its open architecture allows access without the need to alter current technology.

Although specific embodiments have been illustrated and described herein, it will be appreciated by those of ordinary skill in the art that any arrangement which is calculated to achieve the same purpose may be substituted for the specific embodiment shown. This application is intended to cover any adaptations or variations of the present invention. Therefore, it is intended that this invention be limited only by the claims and the equivalents thereof.

What is claimed is:

1. A method of limiting access from an external network to documents stored on an internal network, the method comprising:

building a client list, wherein building a client list includes assigning each client to a role;

building a document list naming documents available to clients assigned to the client's role;

receiving a request for a document stored on the internal network;

associating the request with a client;

determining if the requested document is on the list of documents; and

if the requested document is on the list of documents, fetching the requested document as a proxy and sending the requested document to the client.

2. The method according to claim 1, wherein each document has a unique URL and wherein the document list includes the URL of each available document, wherein the step of determining includes the step of determining if the URL of the requested document is included in the document list.

3. The method according to claim 2, wherein building the document list includes displaying available documents in a tree structure within a graphical user interface, wherein displaying includes querying a document control server to obtain a current version of the document list.

4. The method according to claim 1, wherein associating the request with a client includes the step of authenticating the client.

5. The method according to claim 4, wherein each document has a unique URL and wherein the document list includes the URL of each available document, wherein the step of determining includes the step of determining if the URL of the requested document is included in the document list.

6. The method according to claim 5, wherein building the document list includes displaying available documents in a tree structure within a graphical user interface, wherein

US 6,357,010 B1

17

displaying includes the step of querying a document control server to obtain a current version of the document list.

7. A document control system, including:

an internal network;
an external interface;

a document server connected to the internal network, wherein the document server controls access to a plurality of documents, including a first document; and

a document control server, wherein the document control server receives a document request for the first document, determines a user associated with the document request and authenticates the user, wherein the document control server includes a go list processor for determining if the user has authorization to access said first document and a document processor for reading the first document from the document server, cleaning the first document and forwarding a clean version of said first document to the user.

8. The document control system according to claim 7, wherein the external interface includes a firewall connected to an external network, wherein the document control server communicates to the document server through the firewall.

9. The document control system according to claim 7, wherein the document processor acts as a proxy to hide access to the first document.

10. The document control system according to claim 7, wherein the external interface includes a telephone interface into which a business partner can dial to gain access to the document control server.

11. A document control system, including:

an internal network;
an external interface;

a document server connected to the internal network, wherein the document server controls access to a plurality of documents, including a first document; and

a document control server; and

a data owner interface for building a document list of available documents;

wherein the document control server receives a document request from the external interface for the first document, determines a user associated with the document request and authenticates the user; and

wherein the document control server includes a go list processor for determining, based on the document list, if the user has authorization to access said first document.

12. The document control system according to claim 11, wherein the data owner interface includes a graphical user interface which displays the document list in a tree structure, wherein the graphical user interface queries the document control server to obtain a current version of the document list.

13. The document control system according to claim 11, wherein the document control server further includes a document processor for reading the first document from the document server, cleaning the first document and forwarding a clean version of said first document to the user.

14. The document control system according to claim 13, wherein the document processor acts as a proxy to hide access to the first document.

15. The document control system according to claim 14, wherein the data owner interface includes a graphical user

18

interface which displays the document list in a tree structure, wherein the graphical user interface queries the document control server to obtain a current version of the document list.

16. The document control system according to claim 11, wherein the external interface includes a firewall connected to an external network.

17. The document control system according to claim 11, wherein the external interface includes a telephone interface into which a business partner can dial to gain access to the document control server.

18. The document control system according to claim 8, wherein the document control server is connected to the external network and communicates to the firewall through the external network.

19. The document control system according to claim 8, wherein the document control server is connected to a third network and communicates to the firewall through the third network.

20. The document control system according to claim 7, wherein the document control server is connected to the internal network and wherein the external interface includes a firewall connected to an external network, wherein the firewall is configured to receive the document request and to route the document request to the document control server.

21. The document control system according to claim 7, wherein the document control server includes means for translating links embedded in the first document.

22. The document control system according to claim 7, wherein each user is assigned one or more roles and wherein the go list processor restricts access to documents as function of the role under which the user attempts to access the document.

23. In a system having an internal network and an interface to an external network, a method of handling requests from the external network for documents stored on the internal network, the method comprising:

defining one or more users;

defining documents accessible to the users;

receiving a document request from the external network;

determining a user associated with the document request;

authenticating the user associated with the document request;

determining if the user associated with the document request has permission to access the document requested; and

if the user associated with the document request has permission to access the document requested, retrieving the document requested from the internal network, cleaning the document of embedded links and delivering the document to the user associated with the document request.

24. The method according to claim 23, wherein each document has a unique URL and wherein determining if the user associated with the document request has permission to access the document requested includes:

accessing a document list listing the URL of each available document; and

generating an error message if the document requested is not on the document list.

25. The method according to claim 23, wherein defining documents accessible to the users includes assigning each

US 6,357,010 B1

19

user to one or more roles and limiting access to documents as a function of role and wherein determining if the user associated with the document request has permission to access the document requested includes determining if users in the role associated with the document request have permission to access the document requested.

26. The method according to claim 24, wherein defining documents accessible to the users includes assigning each user to one or more roles and limiting access to documents as a function of role and wherein determining if the user associated with the document request has permission to access the document requested includes determining if users in the role associated with the document request have permission to access the document requested.

27. The method according to claim 23, wherein each document request includes an HTTP header and wherein authenticating the user associated with the document request includes retrieving authentication information from the HTTP header.

28. The method according to claim 23, wherein cleaning the document of embedded links includes looking for a server path link and replacing the server path link with a link to an alias.

29. The method according to claim 23, wherein cleaning the document of embedded links includes looking for an absolute path link, determining if the absolute path link is a link which should be hidden and, if the absolute path link is a link which should be hidden, replacing the absolute path link with a different link.

30. In a system having an internal network and an interface to an external network, a method of handling requests from the external network for documents stored on the internal network, the method comprising:

defining a plurality of users, including a first and a second user;

assigning each user to one or more roles, wherein assigning includes assigning the first user to a first role and the second user to a second role;

defining documents accessible to the users, wherein defining includes limiting access to documents as a function of the roles assigned to the user;

receiving a document request from the external network; determining a user and a role associated with the document request;

authenticating the user associated with the document request;

determining if users in the role associated with the document request have permission to access the document requested; and

if users in the role associated with the document request have permission to access the document requested, retrieving the document requested from the internal network and delivering the document to the user associated with the document request.

31. The method according to claim 30, wherein each document has a unique URL and wherein determining if the user associated with the document request has permission to access the document requested includes:

accessing a document list listing the URL of each available document; and

generating an error message if the document requested is not on the document list.

20

32. The method according to claim 30, wherein retrieving the document requested includes cleaning the document of embedded links.

33. The method according to claim 32, wherein cleaning the document of embedded links includes looking for a server path link and replacing the server path link with a link to an alias.

34. The method according to claim 32, wherein cleaning the document of embedded links includes looking for an absolute path link, determining if the absolute path link is a link which should be hidden and, if the absolute path link is a link which should be hidden, replacing the absolute path link with a different link.

35. A computer-readable medium having program code for limiting access from an external network to documents stored on an internal network, the program code comprising:

program code for building a client list, wherein program code for building a client list includes program code for assigning each client to a role;

program code for building a document list naming documents available to clients assigned to the client's role;

program code for receiving a request for a document stored on the internal network;

program code for associating the request with a client;

program code for determining if the requested document is on the list of documents; and

program code for, if the requested document is on the list of documents, fetching the requested document as a proxy and sending the requested document to the client.

36. A computer-readable medium comprising program code, in a system having an internal network and an interface to an external network, for handling requests from the external network for documents stored on the internal network, the program code comprising:

program code for defining one or more users;

program code for defining documents accessible to the users;

program code for receiving a document request from the external network;

program code for determining a user associated with the document request;

program code for authenticating the user associated with the document request;

program code for determining if the user associated with the document request has permission to access the document requested; and

program code for, if the user associated with the document request has permission to access the document requested, retrieving the document requested from the internal network, cleaning the document of embedded links and delivering the document to the user associated with the document request.

37. A computer-readable medium comprising program code, in a system having an internal network and an interface to an external network, for handling requests from the external network for documents stored on the internal network, the program code comprising:

program code for defining a plurality of users, including a first and a second user;

US 6,357,010 B1

21

program code for assigning each user to one or more
roles, wherein assigning includes assigning the first
user to a first role and the second user to a second role;
program code for defining documents accessible to the
users, wherein defining includes limiting access to
documents as a function of the roles assigned to the
user;
program code for receiving a document request from the
external network;
program code for determining a user and a role associated
with the document request;

22

program code for authenticating the user associated with
the document request;
program code for determining if users in the role associ-
ated with the document request have permission to
access the document requested; and
program code for, if users in the role associated with the
document request have permission to access the docu-
ment requested, retrieving the document requested
from the internal network and delivering the document
to the user associated with the document request.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 6,357,010 B1
DATED : March 12, 2002
INVENTOR(S) : Viets et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Title page,

Item [54], delete "SYSTEM AND METHOD FOR CONTROLLING ACCESS TO DOCUMENTS STORED ON AN INTERNAL NETWORK" and insert -- ACCESS CONTROL TO INTERNAL NETWORK DOCUMENTS WITH CLIENT LIST ASSIGNING CLIENT'S ROLE AND DOCUMENT LIST NAMING DOCUMENTS AVAILABLE TO CLIENTS ASSIGNED TO CLIENT'S ROLE --, therefor.

Column 1,

Line 48, delete "company'firewall" and insert -- company's firewall --, therefor.

Column 2,

Line 7, delete "client'role" and insert -- client's role --, therefor.

Column 4,

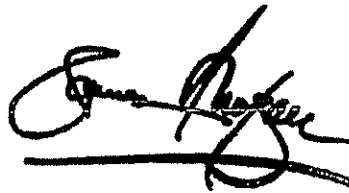
Line 28, delete "roles" and insert -- role --, therefor.

Column 20,

Line 23, delete "client'role" and insert -- client's role --, therefor.

Signed and Sealed this

Twenty-fifth Day of March, 2003

A handwritten signature in black ink, appearing to read "James E. Rogan", written over a horizontal line.

JAMES E. ROGAN
Director of the United States Patent and Trademark Office